H2020-INFRAEDI-2018-2020

EXCELLERAT

The European Centre of Excellence for Engineering Applications

Project Number: 823691

D5.5
Updated Services, Training and Portal Report

Workpackage:	WP5	Centre implementation	
Author(s):	Claudio Arlandini		CINECA
	Eric Pascolo		CINECA
	Lorenzo Zanon		USTUTT
	Francesco Salvadore		CINECA
	Ivan Spisso		CINECA
Approved by	Executive Centre Management		2020-12-31
Reviewer	Michael Wagner		DLR
Reviewer	Niclas Jansson		KTH
Dissemination Level	Public		

Date	Author	Comments	Version	Status
2020-10-05	C. Arlandini	Initial draft	V0.1	Draft
2020-11-25	All	Draft for first review	V0.5	Draft
2020-12-04	C. Arlandini	Clean up after first review	V0.6	Draft
2020-12-09	All	Draft for second review	V0.7	Draft
2020-12-14	C. Arlandini	Final	V1.0	Final

List of abbreviations

ACSI	*American Customer Satisfaction Index*
API	*Application Programming Interface*
CFD	*Computational Fluid Dynamics*
ChEESE	*Centre of Excellence in the domain of Solid Earth*
CMS	*Content Management System*
CoE	*Centre of Excellence*
CSA	*Coordination and Support Action*
DN	*Distinguished Name*
ERD	*Entity-Relation Diagram*
GDPR	*General Data Protection Regulation*
HiFi-Turb	*High-Fidelity LES/DNS Data for Innovative Turbulence Models*
HPC	*High Performance Computing*
ISV	*Independent Software Vendor*
JSON	*JavaScript Object Notation*
KPI	*Key Performance Indicator*
LDAP	*Lightweight Directory Access Protocol*
MS	*Milestone*
ORM	*Object Relational Mapper*
PMO	*Project Management Office*
PRACE	*Partnership for Advanced Computing in Europe*
REST	*Representational state transfer*
RT	*Request Tracker*
SME	*Small Medium Enterprise*
SQL	*Standard Query Language*
SWAN	*System for the Worldwide Exchange of Engineering Data*
WP	*Work Package*

Executive Summary

This deliverable D5.5 "*Updated Services, Training and Portal Report*" presents the outcomes of the second year of Work Package (WP) 5 activities for all its five tasks and in particular for what concerns the architecture and implementation of the service portal. A new version, completely redesigned and with enlarged content, was launched in July 2020 as described in D5.4 "*Updated Portal Release*". Year 2 KPIs are also presented and commented.

After an Introduction section, Section 2 "*Service portal workflows*" presents the main areas of the new release of the service portal, acting as a user guide. This was derived from the effort of Tasks 5.1 and 5.2. Section 3 "*Further applications*" describes Task 5.3 activities related to the onboarding of new applications in the project. Three in particular have been selected, and work is ongoing. Section 4 "*Training*" describes Task 5.4 activities in preparing and deploying an EXCELLERAT training portfolio, while Section 5 "*HPC Provisioning*" accounts for Task 5.5 effort in providing the HPC resources the CoE activities rely on. Section 6 draws a conclusion and outlines future activities.

Table of Contents

Table of Figures

Table of Tables

1 Introduction

The objective of WP5 "Centre Implementation" is ensuring the management of the day-by-day operation of EXCELLERAT as well as supporting the implementation of the defined services. WP5 also manages services, users' access, and related issues; both for internal services and for services towards the users.

This includes:

- Services for end-users as designed in WPs 2,3,4 with a particular focus on the industrial users,
- Training and education services,
- Centre of Excellence (CoE) internal administrative services, and
- Support for new application codes.

Another key aspect is the management of HPC resources provisioning (CPU hours, storage, etc.), including the integration of infrastructure services, for EXCELLERAT activities, to be coordinated with the provision services of the HPC centres.

The main output of this work package is to provide a single-entry point to its ecosystem of services, thus, building up an online access point – hereafter named the Portal (distinct but integrated with the project website, hereafter named Website) – which will be incrementally extended by new services. A new complete and fully redesigned version of the portal was released in M20, significantly before the DoA deadline, as described in D5.4 "*Updated Portal Release*".

This deliverable "*Updated Services, Training and Portal Report*" presents the outcomes of the second year of WP5 activities for all its five tasks and in particular for what concerns the implementation of the above-mentioned portal from a user point of view.

The determination of the KPIs for Year 2 as defined in D5.1 "*Initial Assessment of Training Needs and Services Building Plan*" to support a successful implementation is also presented, together with a workplan for the final year of the project.

2 Service Portal Workflows

The portal offers a wide range of features accessible according to access privileges. The main functions refer to two main areas:

- Content Management System: management of contents relating to EXCELLERAT, i.e. general information relating to project members and infrastructures, and thematic contents authored by the project partners on topics relevant to the Center of Excellence and for possible external users operating on similar issues
- Request Tracker: the portal manages the service requests from the user, and allows the interaction between user and EXCELLERAT staff to be carried out, through dashboards and specific interfaces.

The following describes the workflows for using the portal according to the type of privilege guaranteed by the user.

2.1 Visitor Workflow

The visitor user - not logged in or logged in but who does not belong to any organization - accesses the contents of the portal through two possible paths:

- A usability-oriented path that presents the different types of content within 3 dashboards, i.e. dashboard community, engineer, or developer
- A speed-oriented access path based on the simple list of the different types of content that can be clicked from a classic dropdown menu that can be activated from the header bar.

2.1.1 Dashboards

The dashboards are clearly presented on the Portal landing page (see Figure 1) and are designed to welcome the visitor user, thus framing the different professional profiles of users who can profitably visit the portal. By accessing each dashboard, the user has the opportunity to have a quick look at the types of content most likely suitable for the corresponding profile. For each type of content, the dashboard allows you to have a quick look at a particular content or to press a button to access the list (with search functionality) of all the contents of that type.
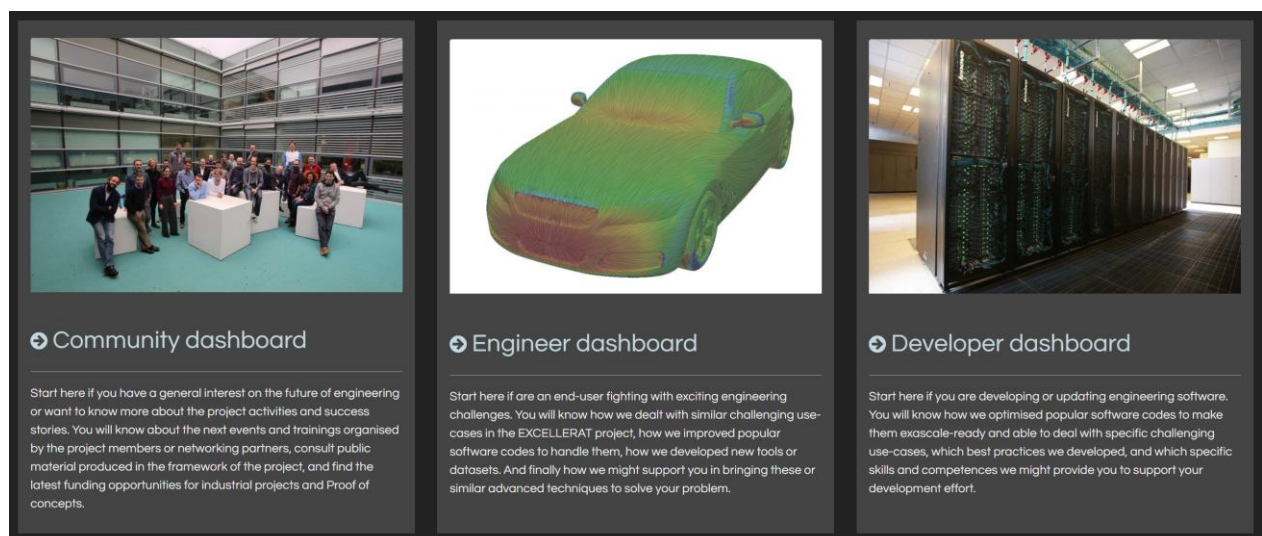
Figure 1: Home page excerpt from EXCELLERAT Service Portal. The three dashboards (Community, Engineer and Developer) are presented and shortly described.

The community dashboard (see Figure 2) focuses on content related to:

- General events such as conferences or project meetings. You can access the list of events with the start dates highlighted. For each event you can access the details page.
- Training activities on topics and skills useful in the fields of application addressed within the CoE. For details on the presentation and user interaction in the context of training activities, please refer to the Section 4 of this deliverable.
- Funding opportunities. For these contents, you access a panel that provides the list of opportunities, highlighting the expiry dates of the calls
- Public dissemination material to be downloaded from the portal itself or linked to other reference sites. For example, journal articles, conference proceedings, HPC center documents, etc. In addition to the list of material, by accessing the page of a certain content, you can download the material directly from the portal or follow an external link to access the material itself
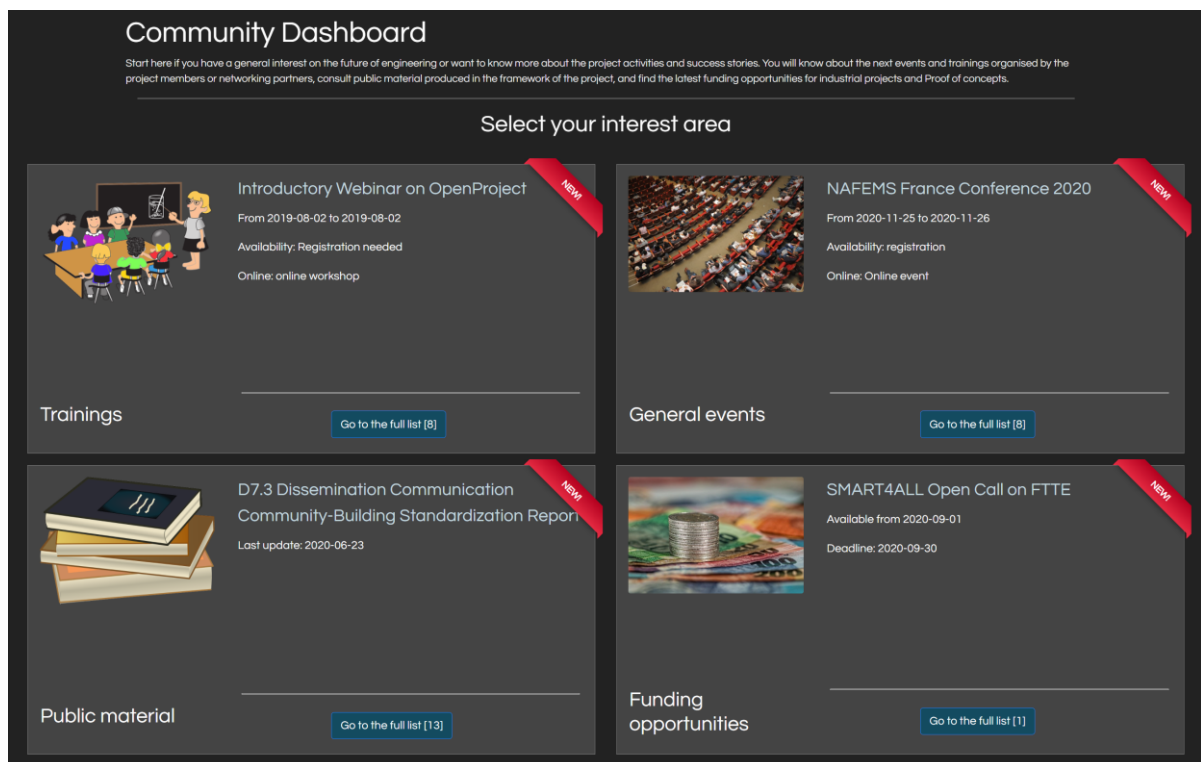
Figure 2: Community dashboard. Trainings, General events, Public material, and Funding opportunities sections are available.

The engineer dashboard (see Figure 3) has 5 sections:

- Industrial use-cases: the section includes concrete cases where the codes and machines of EXCELLERAT are used to get state-of-the-art solutions of typical engineering problems. Use-cases show the power of the HPC resources in a real-world context and may inspire users for their specific application cases.
- Application core codes: these are basic ingredient of the EXCELLERAT projects since they allow for achieving the cutting-edge results of engineering objectives. A number of codes are officially supported within the Services provided in the context of the EXCELLERAT Center of Excellence.

- Tools: are a particular type of Application software supported in the EXCELLERAT context. These are task specific software normally useful in an engineering workflow such as visualization or data transfer software.
- Datasets: these are reference datasets, offered through links to external sites, which can be useful in engineering workflows, such as benchmarks or validation data.
- Training events: same content also present in the community dashboard.
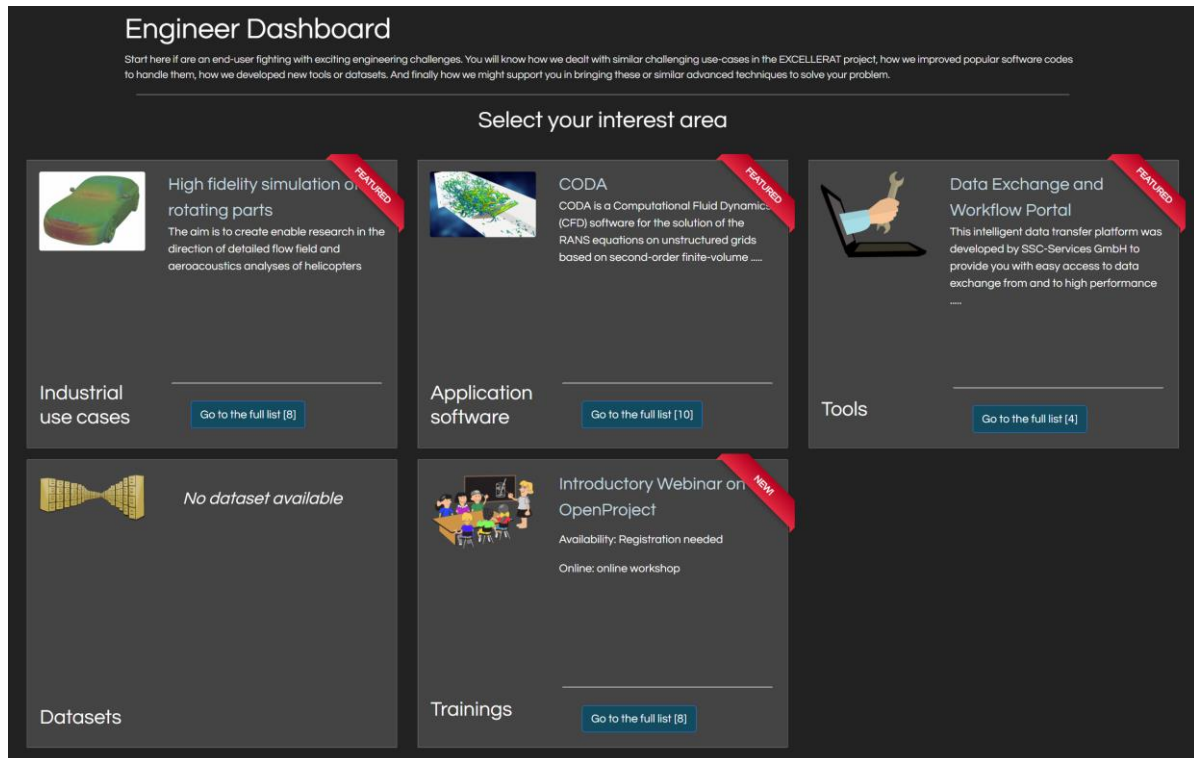
Figure 3: Engineer dashboard. Industrial use-cases, Application software, Tools, Datasets, Trainings sections are available.

The developer dashboard (see Figure 4) has 5 sections:

- Industrial use-cases: same content as the engineer dashboard.
- Application core codes: same content as the engineer dashboard.
- Trainings: same community content and engineer dashboard.
- Best practice: documentary content guides more oriented towards user guides, recommendations for the use of HPC clusters or specific computing architectures.
- Consulting topics: these are consulting services for predefined thematic areas - cataloged according to tags - which help the user to frame their concrete problems under a more general and systematic profile.
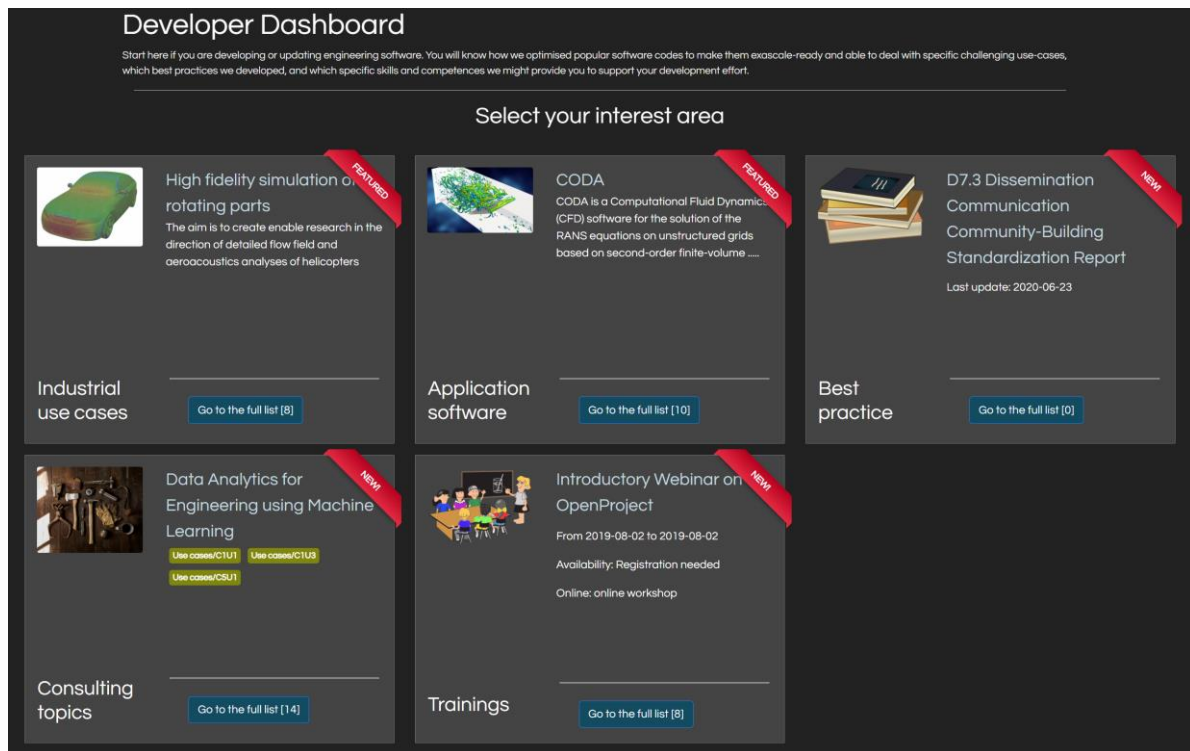
Figure 4: Developer dashboard. Industrial use-cases, Application software, Best practice, Consulting topics, Trainings sections are available.

2.1.2 Header Menu Content Access

In addition to the dashboards, you can access the various content pages following a classic path that starts from the dropdown in the portal header bar. This dropdown shows the available content in a more schematic way, and can be convenient for quick access when the content structure is already clear. Some content, not necessarily essential for the average visitor user, can only be accessed via dropdown and not from dashboards. The dropdown menu has 8 types of content, which mostly follow the types of content defined in the database of the portal backend. In particular, the dropdown allows you to access:

- Consulting topics: contents accessible also from dashboards
- HPC resources: access to the list and descriptions of the computing clusters available among the project members
- Industrial use cases: contents accessible also from dashboards
- Partners: access to the list and descriptions of the 13 EXCELLERAT members
- Interest groups: access to the list and descriptions of institutions / companies that have expressed an explicit interest in EXCELLERAT, accepting to be published on the portal in this context
- Application software: contents also accessible from dashboards. In particular, this content includes 3 sub-types of content, of which the first and third are accessible separately from the dashboards, namely:
  - Application core codes: accessible from dashboards
  - Onboarded codes: these are additional codes of interest for one or more groups of users who have applied for the onboarding procedure to receive specific support within the project. For details on the onboarding procedure, see Section 3.1.3.
  - Tools: accessible from dashboards

- Events. This type of content includes 3 subtypes:
  - General Events: accessible from dashboards
  - Training events: accessible from the dashboards
  - Funding opportunities: accessible from the dashboards. Although they are not really events, from the point of view of the organization of the contents they can be included among the events due to the importance of the temporal aspect.
- Repositories. Content also accessible from dashboards. This content type includes 3 subtypes
  - Public material: accessible from dashboards
  - Best practice: accessible from dashboards
  - Dataset: accessible from dashboards

2.1.3 Content Search and Related Content

For each type of content, from the dashboard to access the page that lists the available content, proposing the most relevant metadata and following the most appropriate display format. From the content list page, it is possible, as an additional function, to carry out free searches using an appropriate search field that matches the title or description fields of the content in question. For example, Figure 5 shows the Industrial use-cases containing the word "aerodynamics".
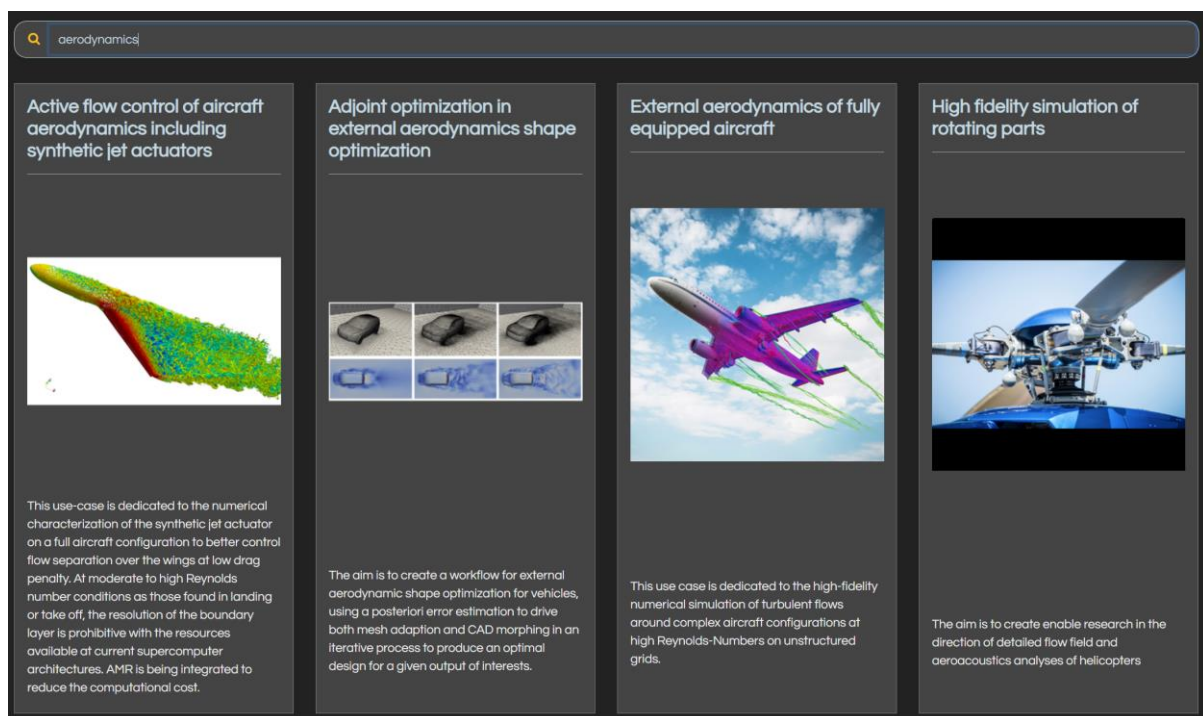
Figure 5: Industrial use-cases search pages. As an example, the results match the correspondence of "aerodynamics" within the description of available use-cases.

Furthermore, once a certain content has been selected and after having accessed the relative page, it is possible to access similar contents - both of the same type and of another type - simply by clicking on the Related Contents proposals that the system processes on the basis of an advanced mechanism of tagging. Thus, a certain industrial use-case can be linked to some codes, to some events, to some trainings and so on. This mechanism is designed to help the user

better understand how to meet their needs. Figure 6 shows a part of the page relating to an industrial use-case with the related contents section highlighted on the right.

Figure 6: Industrial use-case example page. Three boxes are available, namely: Industrial use case information, Request additional support, Related material.

2.2 User Workflow

The logged in user who belongs to at least one organization can benefit from the second main functionality of the portal which is to manage user interaction with the EXCELLERAT staff regarding concrete issues relating to the CoE's own themes.

The EXCELLERAT user-staff interaction revolves around service requests in which users can submit their more or less concrete problems, ranging for example from the request for the activation of a new course to support for the parallelization of a code on architecture hybrid.

In order to make a service request, the user must first create a project. In fact, the service requests are grouped within projects to allow more adequate management in the event that multiple requests are interrelated, as they are linked to a similar concrete objective or even simply temporally close. When creating the project, the user must specify:

- a name
- a description
- which organization, among those to which the user belongs in the portal, should be the owner of the project and, consequently, of the service requests made in that project. In this way, the service requests are always belonging to groups (organizations) rather than to the single user, simplifying the possible handover between users within companies and allowing collaboration between users of the same organization.

The creation of a service request within a certain project is linked to the type of content for which assistance is requested, and is therefore done by clicking on the buttons in the appropriate boxes, visible on the portal in correspondence with the various contents (see Figure 7). For example, by accessing a certain use-case, it is possible to request support on a similar engineering problem, while by accessing the list of events, it is proposed to request a new event to satisfy a certain need.
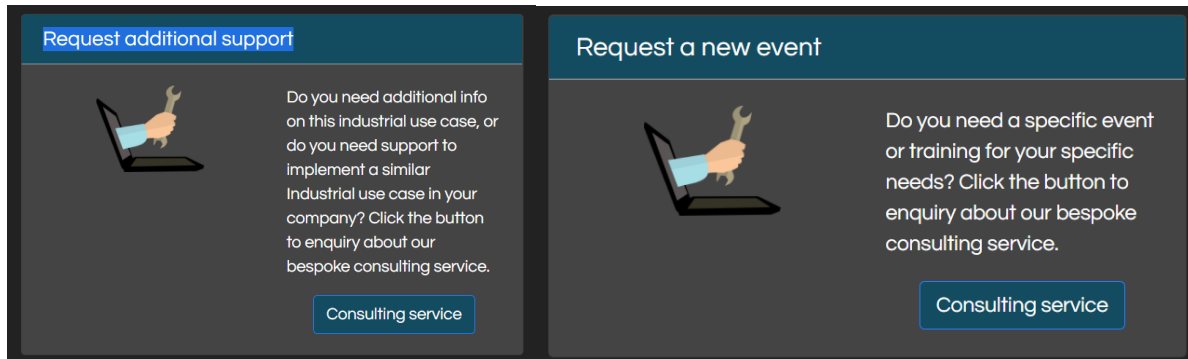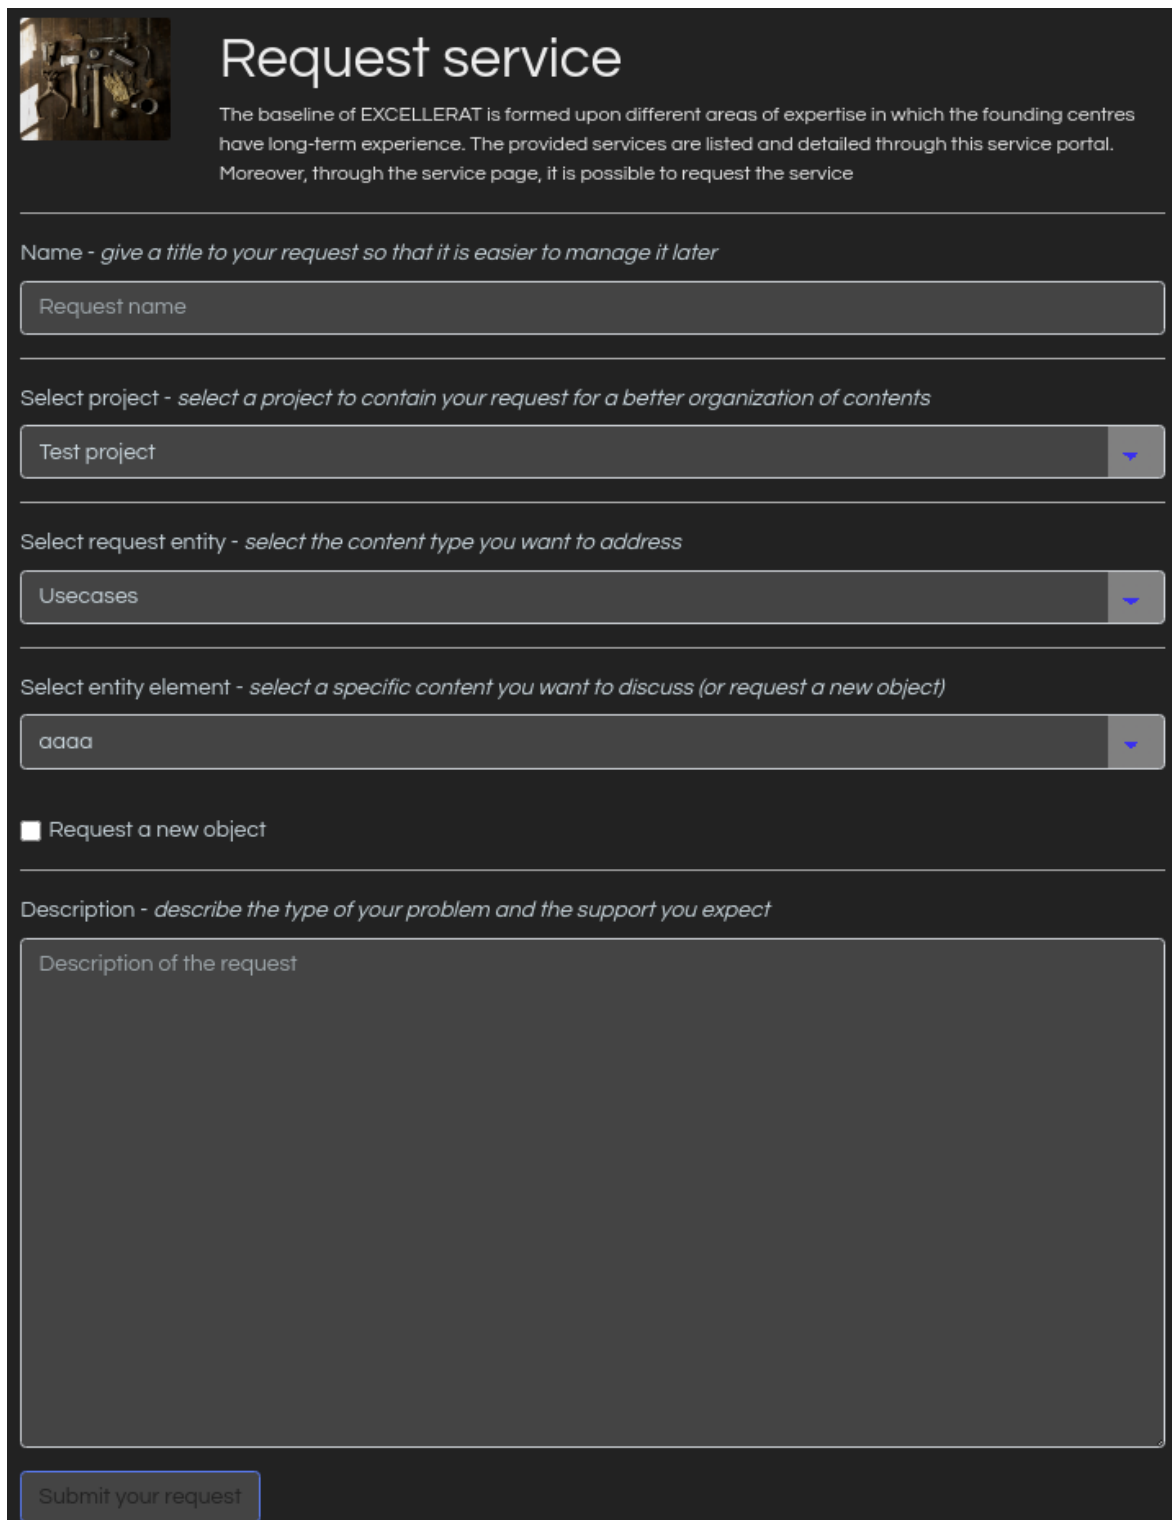
Figure 7: Two examples of boxes containing message and button to apply for a Service Request from EXCELLERAT Service Portal.

By clicking on "Consulting service" you then access the service request page, preconfigured according to the page of origin. The service request page (see Figure 8) requires you to specify:

- an identifying name for the request
- the project to which the user belongs, chosen from the list of projects of the organization to which the user belongs
- the type of content to which the request refers (e.g. Codes or Events) which is proposed pre-filled according to the page of origin
- if the request concerns the creation of a new content (for example a new event or a new code) or if, on the other hand, support is requested for some content already present in the portal
- a free description of the support request, in which the user can describe the type of request with a first level of detail so that the EXCELLERAT staff can consider whether it is a type of request that the CoE can consider.

Figure 8: Service Request web form. The proposed fields have to be filled to submit a valid request for service.

After opening the request, you get to the page for viewing and developing the service request that manages the interaction with the EXCELLERAT staff. The interaction can take place in 2 ways:

- exchange of messages: the user or other members of the organization to which the project belongs can write messages, the EXCELLERAT staff can reply to messages or

can also exchange private messages that remain visible only to other staff members of the project
- file exchange: there is a reserved area to manage files shared between user and staff, which can be used simply through the web page and can be used to exchange documents or even files of codes and configurations relevant to the service request in question.

Each service request has a status that can only be changed by staff users within the created / defined / running / completed / canceled set.

The projects of the organizations to which a certain user belongs are accessible from the specific "projects" dashboard from which you can see the essential information for each project, namely:

- name
- creation date
- owner organization
- user creator
- number of associated requests
- button to access the project

By accessing the project page, you access the project dashboard (see Figure 9) which shows, alongside the general information on the project, the list of service requests relating to the project itself in which the following are specified:

- name
- creation time
- creator user
- entity, or the abbreviation of the type of content for which the request was activated, such as "codes" or "trainings"
- status
- the number of messages exchanged in the service request
- button to access the service request

Figure 9: Example of a project page. Information (left) and service requests dashboard (right) are shown.

By accessing the single service request, you are taken to the service request management page described above which allows user / staff interaction.

2.3 Staff Workflow

Portal users who belong to organizations flagged as "staff" are EXCELLERAT staff users and can access special sections of the portal. In particular:

- can access the EXCELLERAT internal services page
- they can interact in the context of service requests opened by external users with the role of staff users who respond to user requests

2.3.1 Internal Services

Figure 10 shows the Internal Services page from which you can quickly access the main six productivity tools for CoE members.
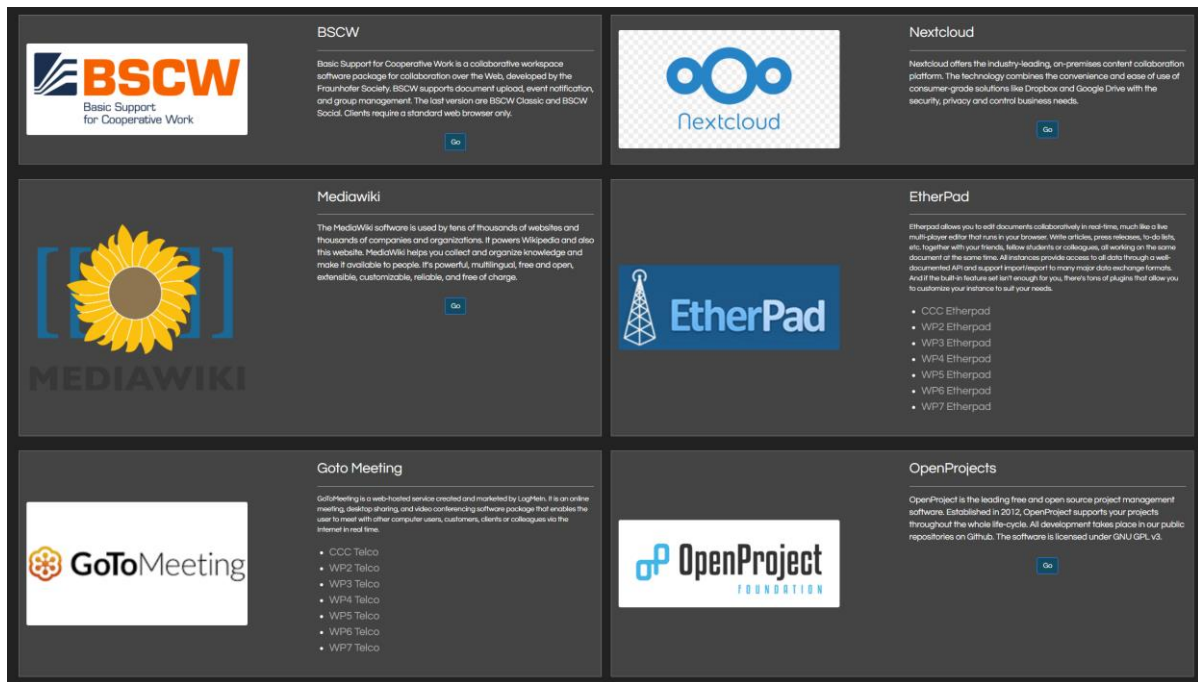
Figure 10: Internal services page in EXCELLERAT Service Portal. Six services are available through different platforms, i.e. BSCW, Nextcloud, MediaWiki, EtherPad, GoToMeeting, OpenProject.

In particular:

- BSCW [1] as document repository and collaborative workspace software package for collaboration over the Web (BSCW supports document upload, event notification, and group management)
- Nextcloud [2] as a content collaboration platform.
- MediaWiki [3] is a collaboration and documentation platform brought to you by a vibrant community
- EtherPad [4] for collaborative editing of documents
- GoToMeeting [5] for online meetings
- OpenProject [6] as project management software

These tools were implemented thanks to the ongoing analysis of Task 5.1 on needed internal services. The solutions adopted are open-source solutions except GoToMeeting. The choice of open-source has allowed, in addition to eliminating the license costs, to use installations on servers of the EXCELLERAT centers, thus ensuring security and full control of the data involved.

2.3.2 Request Management

Staff users can access the service request management panel from the header bar. From the panel it is possible to see the basic information of the service requests, so that the staff user can decide to investigate the type of request and possibly respond to the user who made the request. By accessing the request, you get to the same page seen by the user but with special privileges that allow you to:

- Change the status of the request
- Interact with staff messages for the user
- Interact with comments type messages visible only to other staff users
- Interact by exchanging files (similarly to what the user can do).

2.4 Admin Workflow

Users belonging to the particular admin group have the most powerful privileges in the portal. In particular, the admin user is responsible for:

- Management of organizations
- Creation of new content in the portal

Organizational management includes the creation / editing of the organization's metadata and the management of users who belong to it. The page for creating organizations (see Figure 11) requires you to indicate:

- name
- description
- image (as logo)
- "staff" flag to be activated if necessary
- list of users who belong to the organization

Figure 11: Admin web form to create a new organization, i.e. a group of users which share ownership of requests.

It is then possible to access the list of organizations or view the characteristics of a certain organization.

As for the creation of portal content, the users of the admin group are the only ones authorized to create new content, which is done by filling in forms that require the characteristics to be proposed in a different way for each type of content. In order to create a coherent method of content creation by different authors - ensuring a good uniformity of quality and method of presentation of the contents themselves - the creation of a content for the portal is managed by the members of WP5 in coordination with the authors of the different pages through a standardized procedure. The procedure goes through the compilation of a .docx template corresponding to the type of content to be created. The template actually contains the same fields that will then be inserted in the creation web form. However, the intermediate step of the template document allows on the one hand the authors to dedicate themselves completely to the contents avoiding having to interact directly with the portal, on the other hand it gives the members of WP5 EXCELLERAT the possibility to verify the written contents with respect to the uniformity criteria and established qualities for each type of content. In addition, it is possible to analyze the free tags entered by the author of the content and understand if there are any other existing tags that can be assimilated, and if the proposed tag is actually useful for use

in a web application to refine or link content together. This tagging work makes the intermediate step through templates always highly recommendable. However, for the creation of some types of content (especially training and events) some members of EXCELLERAT can proceed directly with the insertion of the content in the portal, avoiding the intermediate step of the template for reasons of efficiency.

2.5 KPI Analysis

Concerning the KPIs collection, we will start from those defined in D5.1. In that report, a distinction was made between the KPIs for internal services, defined by Task 5.1, and those for external services, defined by Task T5.2.

KPI number	Description	Definition	Target	Score at M24
5.1.1	Percentage of the number of internal services activated compared to the number of services identified for the present implementation release.	*N_activated / N_identified * 100*	>75%	92% (6,5 implemented / 7 high priority)
5.1.2	Number of internal active users	*N_active_monthly*	≥15 for the first year after MS4 completion, with a 20% average increase later	20 (estimated by telco and BSCW interaction)
5.1.3	ACSI score: user satisfaction for internal services, annually	$N_{ACSI\_int\_services}$	≥80	66

Table 1: KPIs for Task 5.1

Table 1 gives an overview of the KPIs for internal services.

As regards the ACSI score, it was measured thanks to an online survey. The evaluation, although positive, still does not reach the desired target. Analyzing the answers provided by the participants, we find relevant and interesting observations. In particular, most notes are specific for some tools:

- OpenProject is not enough user-friendly
- Etherpad does not work sufficiently well for collaborative editing
- MediaWiki would be easier to use with the WYSIWYG extension

While other notes are more general:

- Tools require more integration among each other
- Guidelines are missing so as to keep things tidy and uniform and allow for easier search and find

In the continuation of the Project, it will be evaluated if and how to address these points to improve the experience of using internal services.

KPI number	Description	Definition	Target	Score at M24
5.2.1	Number of external	*N_ext_users*	≥20 at the end of the first year after MS4 completion, ≥30 at the end of the project	N/A
5.2.2	Number of requested services (for the present implementation release)	*N_services*	>1	N/A
5.2.3	Percentage of the number of external services activated compared to the number of services identified for the present implementation release.	*N_activated / N_identified * 100*	>75%	N/A
5.2.4	Uptime	*Time of portal availability /Time from first production release*100*	>70% for the first year after MS4 completion >85% for the second, >95% later	N/A
5.2.5	ACSI score: user satisfaction for external service, annually	$N_{ACSI\_ext\_services}$	≥80	N/A

Table 2: KPIs for Task 5.2

Table 2 gives an overview of the KPIs for external services. During the second year of the project, the portal has been thoroughly revisited and put into production and, therefore, also available for external users. However, we think metrics relating to the availability of the portal and the access and use of external users are not yet significant since the time between the opening and the writing of this deliverable has been too short. For this reason, please refer to the next iteration of this deliverable for the evaluation of the metrics of this table.

2.6 Perspectives

The portal is, as mentioned, currently in production. However, there are several aspects on which it is appropriate to continue development to improve the current functionalities and/or add new ones.

In this context, we just want to highlight two points that are directly related to the user experience:

- The final version of the portal will contain a new dashboard, for visitors belonging to a new user segment, HPC vendors companies. This activity follows the directives on the segment definition, and the relative services appealing to it, published in Deliverable D4.2.
- The development of the Portal and the Website has evolved in parallel ways. While the two remain distinct, it is now time to better integrate the two efforts, in terms of avoiding duplicate content and a similar look-and-feel, for the final version of the Portal at Year 3. A taskforce combining effort of WP5 and WP1 is working on that and the results should be available by January 2021.

3 Further Applications

3.1 Status and Activities Performed

In D5.1 *"Initial Assessment of Training Needs and Services Building Plan"*, several ways have been foreseen for EXCELLERAT to attract new applications and, thus, potential new end-users and customers. Among the main ones, the progress in the interaction with the Interest Groups since D5.3 *"Services, Training and Portal Report"* is described in detail in Section 3.1.1. Another one regards the Portal development, object of Section 2 of this document and of D5.4 *"Updated Portal Release"*. The particular issue of the onboarding workflow in the Portal is described in Section 3.1.3 below.

Other ways of attracting applications involve the outreach of EXCELLERAT members and the recruitment through awareness creation channels and contacts at (virtual) conferences and workshops. It is the case of the applications PAinG-Flow and HiDALGO-UAP, about which more details are given in Section 3.1.2.

The onboarding of further codes has technical and scientific advantages for both EXCELLERAT and the external codes involved. These can benefit from the EXCELLERAT achievements, and at the same time act as proof-of-concept of the transferability of the centre's results. At this stage of the project, the onboarding is also important for multiple other reasons: Testing the procedures for onboarding through the Portal, creating success stories for outreach and dissemination, and showing that there is indeed a need of EXCELLERAT services from the community (therefore creating an exploitation channel).

3.1.1 Interest Groups

The EXCELLERAT Interest Groups are recruited to communicate, monitor, and validate the project's overall goals and its industrial, as well as technological relevance for the user communities.

The first contact through "letters of support" marked the kick-off of the interaction with the Interest Groups (D5.1). Next, with an official invitation, the Interest Groups have been asked to fill in and sign a consent form, requesting information about their agreement to become member of the EXCELLERAT Interest Groups, to appear in EXCELLERAT dissemination material, and to subscribe to the Interest Group mailing list (each a separate choice).

The Interest Groups have been so far involved in two virtual meetings. A first meeting hosted by HLRS on November 8, 2019 was dedicated to involve the Interest Groups in the EXCELLERAT activity. A short overview of the project has been given, and three applications have been presented by the respective developers: the data exchange platform SWAN, visualisation with Vistle, and the BigWhoop compression library. During a second web-seminar on June 26, 2020, two new Interest Groups have been introduced to the CoE activities and the refurbished Portal has been displayed.

Those Interest Groups that consented to be publicly mentioned are listed in Table 3 (updated since D5.3), together with their attendance to the virtual meetings. These Interest Groups also appear in both the EXCELLERAT Website (Figure 12) and Portal [7]], with a short description and tags connecting their areas of interest to the multiple project's domains. As an update to D5.3, Leonardo, NVIDIA, and Steinbeis Consulting Center AI (STAI) joined as

EXCELLERAT Interest Groups in 2020. In particular, Leonardo, the largest Italian manufacturing company in fields such as aerospace and defence, has connections to CINECA.

Figure 12: Interest Groups in the EXCELLERAT Website [7].

We can mention at least two ongoing activities together with the Interest Groups, while further plans are described in Section 3.2.1. The interaction with Politecnico di Milano is mainly connected to the Rotorcraft Aerodynamic Lab led by Prof. Luigi Vigevano, whose test-case ROSITA is a candidate for onboarding (see Section 3.1.2). Wim Slagter, Director for High Performance Computing (HPC) and Cloud Alliances at Ansys, was interviewed for the third episode of the EXCELLERAT HPC Podcast. The podcast [8] was broadcast in August 2020 and focused on autonomous vehicles and driving. The interview revolved specifically on the question of how (exascale) simulations could support the development of autonomous vehicles.

Interest Group	Typology	Meeting 1	Meeting 2
ANSYS Germany GmBH	Code-Developer/ISV		
Dassault Aviation	Code-Developer/ISV		
ESI Group	Code-Developer/ISV		
Fraunhofer IGD	Scientific Expert		
Politecnico di Milano	Scientific Expert	Yes	Yes
STAI	Scientific Expert		
University of Ljubljiana	Scientific Expert		
University of Rome Tor Vergata	Scientific Expert	Yes	
AMD	Technology Provider	Yes	
ARM	Technology Provider	Yes	
ATOS	Technology Provider	Yes	Yes
CRAY	Technology Provider	Yes	Yes
HPE	Technology Provider		
Intel	Technology Provider	Yes	
Leonardo	Technology Provider		Yes
NVIDIA	Technology Provider		Yes
Westinghouse	Technology Provider	Yes	

Table 3: Interest Groups and their attendance to the virtual meetings.

3.1.2 Onboarding Applications

A full onboarding requires the completion of the Portal workflow steps, which is still an ongoing work, as described in the next section. Even so, in Y2 and Y3 there have already been contacts and interactions with external applications. Some or all of them will be invited to go through the workflow once definitely approved and implemented.

Two points have been deemed as essential to initiate an onboarding. On the one hand, the application must be consistent with the scope of the project, namely it is an engineering code with potential for high performance and scalability (possibly exascaling). This must be confirmed by the Project Management Office (PMO). On the other hand, at least one EXCELLERAT partner must have available effort within the project (or as in-kind contribution) to work on it. So far, three applications fulfilled these constraints.

PAinG-Flow aims at exerting the particle-fluid momentum coupling in multiphase turbulent flows on a code able to run on supercomputers with GPUs (Figure 13). Such a code will solve the incompressible Navier-Stokes equations in wall-bounded and free-shear conditions. The team FLUMACS (FLUid dynamics of MAcro and micro Complex Systems) at the Mechanical and Aerospace Department at the Sapienza University of Rome requested EXCELLERAT support of CINECA HPC specialists. This collaboration has been accepted by the PMO in February 2020, hence a workplan has been elaborated, which can be found in the Annex 1:

EXCELLERAT Onboarding Proposal: PAinG-Flow. This includes a thorough description of the project, a time plan, the personnel involved, and the required compute resources.
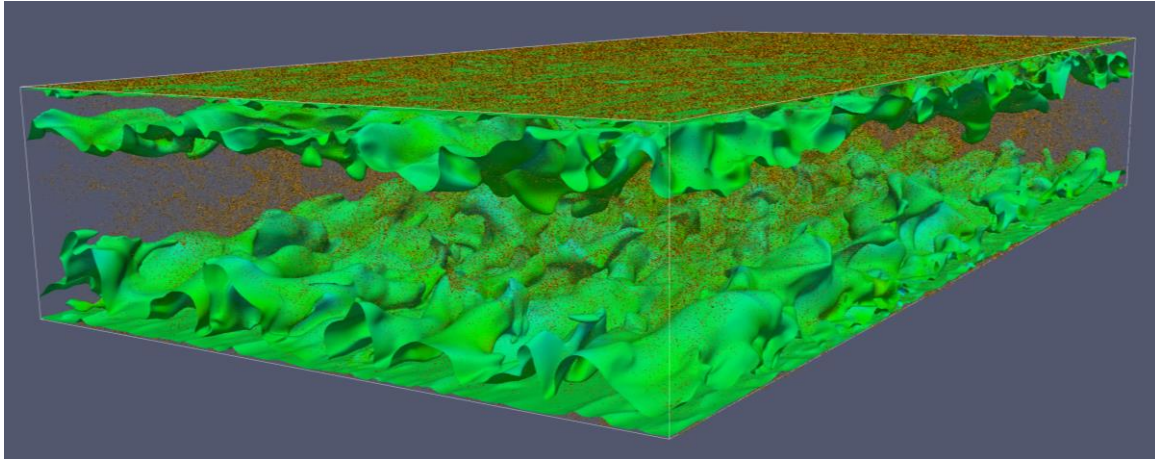
Figure 13: Direct Numerical Simulation of particle-laden turbulent channel flow using Exact Regularized Point Particle (ERPP) method of PAinG-Flow, from [9].

ROSITA (ROtorcraft Software ITAly) [10] is a flow solver which implements the "Chimera approach" [11], capable of performing unsteady Reynolds-averaged Navier-Stokes computations over complex configurations. ROSITA's goal is in particular the simulation of arbitrarily large displacements occurring in rotorcraft aerodynamic simulations (e.g. of helicopters, Figure 14).

The project is led by Prof. Luigi Vigevano with his collaborators at the Aerospace Engineering Department of Politecnico di Milano and counts already an industrial partnership with Agusta Westland. Focusing on this project, Politecnico di Milano has also become an EXCELLERAT Interest Group. The workplan of the activity is under development.
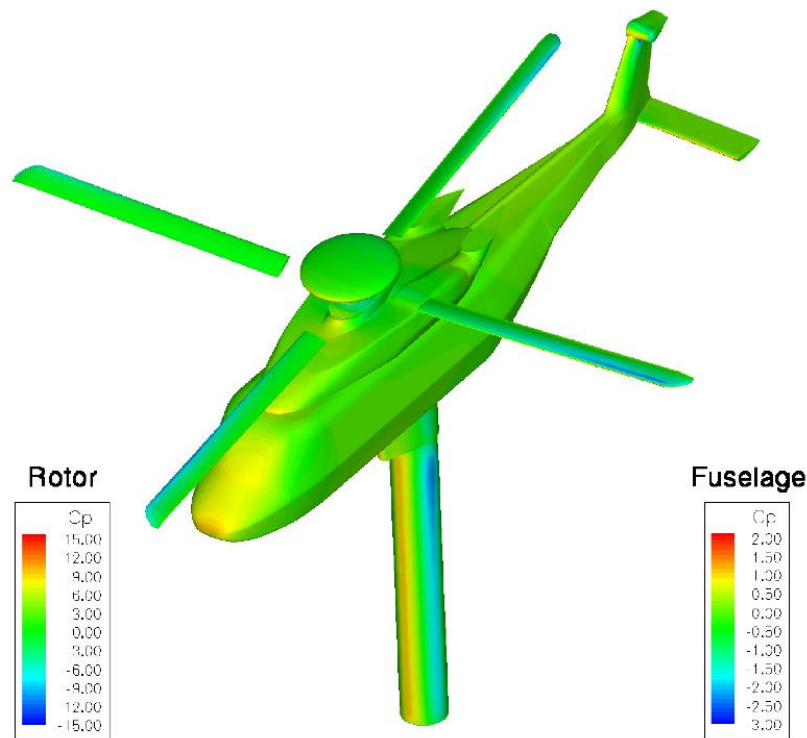
Figure 14: ROSITA: Effects of pylons and walls on the measurements, pressure [10].

HiDALGO-UAP is the third candidate for onboarding originating from the partner CoE HiDALGO [12], see also [13], within the Urban Air Pollution (UAP) Pilot led by the Szechenyi Istvan University (SZE). So far, simulations of the urban air flow have been run with OpenFOAM (Figure 15), but the EXCELLERAT software

In this HiDALGO use-case, the urban air flow is simulated with OpenFOAM (Figure 15).

Results have been presented at the HiPEAC conference in January 2020 [14], where a first contact with EXCELLERAT has been established. Collaboration plans involve the evaluation of the HiDALGO CFD codes by EXCELLERAT, the definition of optimal simulation parameters, the examination and improvement of the parallelisation, and optimisation at different levels (pre-processing, geometry preparation, and mesh generation). As a preliminary result, consultation with EXCELLERAT experts at USTUTT allowed the UAP code to circumvent bottlenecks caused by input-output and to optimise the numerical schemes. The scalability of the UAP simulation consequently improved by about 30%, still in the OpenFOAM framework (June 2020, see [15] for details).

It is envisaged to replicate these simulations on the EXCELLERAT flagship code Nek5000, which would optimally match the HiDALGO requirements of solving incompressible and low Mach number Navier-Stokes equations. Even so, a major difficulty lies in the complex mesh, which would require a significant development effort of the currently available tools.
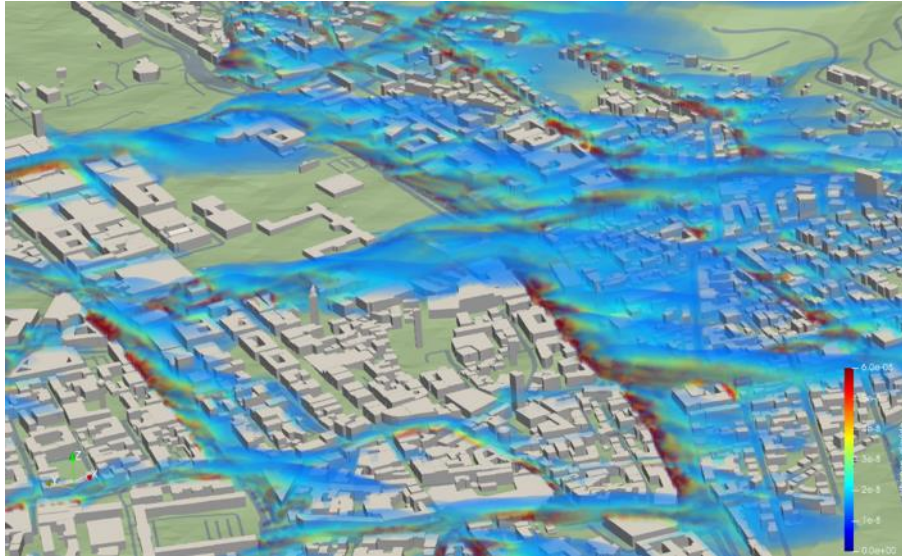
Figure 15: HiDALGO-UAP [12]: Simulated NO_2 concentrations in Stuttgart.

3.1.3 EXCELLERAT Portal Onboarding Workflow

At the beginning of the project, the structural requirements for the Portal in order to attract external codes have been listed. Some of them had already been fulfilled in the first provisional Portal release (see D5.3) and therefore will not be discussed again. Further ones have been implemented in the current second release:

- *A description of the possible target audience*: The Portal has three dashboards displaying content targeted respectively to the engineering, developers' and public community. The "Consulting Topics" page is in particular devoted to external engineering applications of different fields that could benefit from EXCELLERAT achievements by requesting a service.

- *The requirements they [the target audience] have to fulfil*: In this initial period of availability of the onboarding process, we do not expect a number of applications exceeding our review capabilities. Therefore, the selection of successful external applications and codes is delegated to the contact persons of the requested service. It is planned to fill in the page [16] according to the criteria listed in D5.1 (Section 4.2.3).

- *A clear description of the services that can be provided* is provided in the "Consulting Topics" section of the Portal.

- *A direct link to a simple contact form*: Each page of the Portal displays an email address enabling (new) users to directly contact the EXCELLERAT consortium[1]. The access request by email will be soon replaced by a contact form (see Section 4.1 of D5.4).

- *A direct link to the service/support request form*: Once a Portal account has been activated for a new user, the logged-in applicant who would like to select a service is directed to the creation of a project. When requesting a service, the user should provide a description of their issue as in the onboarding workflow described below.

- *Documentation to guide interested parties through the sections of the portal*: Every section of the Portal is introduced by a header describing the underlying page functions. Moreover, this public deliverable contains in Section 2 a thorough explanation of the Portal features.

- *Information about Interest Group members* have been added separately from the Partners' ones, as planned in D5.3 and described in Section 3.1.1.

---

[1] *support@excellerat.eu*

- *Further material that is helpful for engineering applications:* All sections of the Portal (Consulting services, Application software, Industrial use cases, etc.) are continuously updated according to the evolution of the flagship applications.

Based on the assessment of portals offering similar services, the necessary mechanisms which will enable external entities to apply for support for their codes have been elaborated and graphically diagrammed with Bizagi modeler [17]. This diagram contained in detail the various aspects of the onboarding of a new application, also distinguishing among the different responsible actors and their interaction (EXCELLERAT Portal, panel, service-provider, administration). A compromise between this initial idea and a feasible workflow compatible with the evolution of the Portal had to be found and led to the Onboarding workflow below. The implementation is a work in progress and will require the approval by the PMO before public release.

After successful registration to the Portal, the user creates a project by compiling and submitting a provided form. This form, under implementation, contains all fields necessary to get a clear view of the kind of request: general description of the application, technical and scientific domain, methodology so far applied, envisaged or ongoing applications, current scalability performance, already detected bottlenecks, whether the application is fully or partially under NDA, etc. The form should be thorough enough to avoid too high a number of further inquiries. For this reason, it will be elaborated by WP5 together with the technical WPs and, if necessary, the PMO.

After submission, the second step is carried out by an evaluation team composed by leaders of WP 2 *"Application Development"*, 3 *"Driving Exa-HPC Methodologies and Technologies"*, 4 *"Enhanced services"*, 5 *"Centre Implementation"*, and a PMO representative. General evaluation criteria include a feasibility checklist to pre-screen the application, and at least one partner having effort and interest to work on the submitted code. The evaluation might involve the request of additional information via tele-conferences or emails, and should be concluded within a specified timeframe. All communication with the applicant will happen through the "Service requests" feature [18], which functions as a request tracker (see Section 2).

In case of positive evaluation, both the partner selected for collaboration and the applicant will receive a message on their respective Portal accounts. In particular, the applicant will receive a template for their "onboarded application software" to provide the information to be displayed on the Portal [19]. The applicant will also decide which results and advances they would like to post on the EXCELLERAT Portal during the joint work.

Within a specified timeframe, the designated EXCELLERAT partner will produce a workplan with a time-schedule and submit it to the applicant. Finally, the applicant will be also requested to evaluate the onboarding procedure (KPI 5.3.3) through an already available survey [20].

In case of negative evaluation, the applicant will receive a message on their Portal account, containing an explanation of the rejection. The applicant will be requested to evaluate the onboarding procedure (KPI 5.3.3) through the survey mentioned above.

3.2 Activities Planned

3.2.1 Interest Groups

New Interest Groups might join in Y3 of the project, and a new virtual meeting will be scheduled. They will be more involved in the EXCELLERAT outreach, e.g. by adding "exclusive" Interest Groups content to the newsletter, and involving them further in

contributing to the podcast (see also D7.4 *"Updated Dissemination, Communication, Collaboration, Community Building and Standardization Report"*). Some of the training activities will be especially directed to the Interest Groups, such as the "Data Analytics" workshop (Section 4.3).

They will be encouraged to actively collaborate to the technical aspects of the project, especially whenever they can contribute with promising applications and codes. The planned collaboration on a test case present in both EXCELLERAT and exaFOAM (see Section 3.2.2) will in fact happen in the frame of the Interest Groups, since ESI Group is the Principal Investigator of exaFOAM. AMD's partnership with USTUTT will be strengthened by the workshop planned (virtually) at USTUTT in January 2021 (see Section 4.3).

3.2.2 EXCELLERAT Onboarding Development

An interesting ongoing interaction is that of EXCELLERAT with OpenFOAM. This takes place mainly during the OpenFOAM Thematic Working Group (TWG) bi-monthly conference calls, whose purpose is to exchange information and best practices regarding the use of OpenFOAM in the context of an HPC-Cloud. The group web-page is hosted on the EXCELLERAT Wiki. In order to achieve a concrete collaboration, OpenFOAM members presented challenges and open problems that could potentially overlap with and find support by the EXCELLERAT codes and tools.

An interesting test case was raised by members of the University of Minho, Portugal, working on complex fluids (polymers in viscoelastic regime) in the frame of structural mechanics and fluid-structure interaction (see [21]). EXCELLERAT support provided by FEniCS and the Data Transfer Platform has been evaluated but eventually deemed not successful by the PMO.

A different kind of collaboration has been envisaged within the EuroHPC project exaFOAM, due to start in Q1/2021, and to which most members of the TWG will be affiliated. Moreover, ESI Group is the Principal Investigator of exaFOAM and is already an EXCELLERAT Interest Group. exaFOAM will work with industrial partners on DrivAer [22], which has been chosen as a benchmark to be presented at the "2nd Automotive CFD Prediction Workshop" in August 2021 (see [23] for the first edition). Since DrivAer is at the same time one of FEniCS' test cases, a collaboration with exaFOAM has been envisaged. The technical work will begin at exaFOAM's start, but preliminary discussions are already planned during the upcoming TWG calls.

Awareness creation channels and presentations by EXCELLERAT members at (virtual) conferences and workshops should foster the onboarding of further applications and the dissemination of their evolution story. The WP5 cooperation with WP7 *"Awareness, Impact Creation and Outreach"* has been strengthened since D5.3. On the one hand, T5.3 *"Further Applications"* and T7.4 *"Community building"* are monitoring together the interaction between EXCELLERAT and OpenFOAM. On the other hand, new codes and the Interest Groups will be contacted by T7.1 *"Dissemination and communication"* to contribute to the EXCELLERAT newsletter and podcast (see also D7.4).

3.3 KPI Analysis

In D5.1, four functional and quality-related KPIs have been proposed, which can be seen in Table 4. The quality-related KPIs refer to the ACSI Score (American Customer Satisfaction Index), defined in Section 2.3 of D5.1.

As explained in D5.3, KPIs for this task could be measured only after the first Portal release (MS4 in M14). However, the onboarding workflow has been finalised (subject to PMO

approval) at the second release, concurrent to this deliverable. Nevertheless, we can count three submitted applications (PAinG-Flow, ROSITA, and HiDALGO-UAP) between M15 and M23, corresponding to KPI 5.3.1 score 1 (target 1 reached). Considering the single approved application PAinG-Flow, KPI 5.3.2 reaches score 1/3 (target 1 not reached).

Once the workflow will have been approved and implemented, we expect higher scores for both KPIs 5.3.1 and 5.3.2, as well as the possibility of measuring KPIs 5.3.3 and 5.3.4. In this regard, satisfaction surveys related to both the onboarding and the EXCELLERAT support process have been implemented and can be previewed at [24] and [25].

KPI number	Description	Definition	Target	Score at M23
5.3.1	Number of external entities sending an application (quarterly)	$N\_app\_sent\_quarterly$	≥1	1
5.3.2	Number of selected applications to interact with (quarterly)	$N\_app\_interacting\_quarterly$	≥1	1/3
5.3.3	ACSI score: Satisfaction of the on-boarding (candidature) process for a service (e.g. via survey/feedback sheet) as average per year	$N_{ACSI}\_onboarding$	≥80	N/A
5.3.4	ACSI score: Satisfaction of the processed applicants (e.g. via survey / feedback sheet) as average per year	$N_{ACSI}\_applications\_per\_year$	≥80	N/A

Table 4: KPIs for Task 5.3

4 Training

Due to the COVID-19 pandemic, 2020 has been a very peculiar year in (re-)organising training events, which suffered from gathering restrictions and bans. Even if the obstacles in knowledge delivery could mostly be circumvented via online events, the quality of networking was drastically reduced. Together with WP7, improvements are already being sought for, e.g. by increasing the EXCELLERAT presence in social media (see D7.4).

General and event-specific encountered difficulties and proposed solutions for "teaching during pandemic" are detailed throughout this section.

4.1 Status of Assessment of Training Activities

The assessment of training activities carried out in January 2019, and thoroughly described in D5.1, encompassed events organized, held, or advertised by the organisations participating as members of EXCELLERAT. A regular assessment has the goal of maintaining an updated map of the existing initiatives in training and education, in order to also highlight the emerging needs for structuring the EXCELLERAT training roadmap and initiatives.

After a WP5-internal discussion, this initially biannual assessment has been rescheduled as annual, since the majority of partners plan most of their activities with a yearly cadence. Moreover, similar assessment activities are carried out by FocusCoE and also by the Coordination and Support Action for the European National Competence Centres (CASTIEL [26]). At the time of writing, CASTIEL is conducting a thorough training assessment at national level among 33 European countries, which might even replace the EXCELLERAT one in Y3.

The outcome of the EXCELLERAT assessment in Y2 has been presented at the First EXCELLERAT Review in July 2020 and can be summarised as follows. The survey aimed at mapping the existing training initiatives of EXCELLERAT partners, about 60% of which offer training. The responses have been collected between February and April 2020: Due to the concurrent COVID-19 pandemic outbreak, many events had to be rescheduled, offered online, or cancelled. However, the original responses were analysed.

The table in **Errore. L'origine riferimento non è stata trovata.** compares the 2019 and 2020 assessment results, showing a very similar outcome. In 2019, 135 events took place (ranging from half-day- to ten day-events: average 2.8 days), resulting in 382.5 total course days. In 2020, 134 events took place (unchanged duration average 2.8 days), resulting in 378 total course days.

In 2020, there has been a consistent decline in the offer "Parallel Programming" and "Performance Optimisation and Debugging", balanced by a shoot-up in courses of type "Data in HPC". This trend reflects the increasing interest in disciplines such as High-Performance Data Analysis (HPDA) and newer perspectives, such as quantum computing. The number of programming language courses also increased: e.g., USTUTT strengthened their C++ offer, Python courses are spreading due to its affinity to Machine and Deep Learning applications.

The assessment outcome can be compared with the gaps in the offer emerged at the Focus CoE Stakeholder Workshop at the European Commission (Bruxelles, October 8, 2019). There, EXCELLERAT has been represented along with nine other CoEs (see also D5.3 and especially FocusCoE [27] D4.3 *"Report on the outcome of the comprehensive European stakeholder workshop on HPC training and skills development needs"*). A demand for HPC tools in universities already from an early stage was highlighted, reflected in the increase in

programming language courses. Then, the request of large companies and small ones (SMEs) of resources and skills for HPDA coincides with the increase of courses "Data in HPC". Finally, the need of special tools, interfaces, and environments for SMEs does actually not reflect the decrease of courses of type "Special Communities" and "Compute Cluster".

To better support the assessment with an EXCELLERAT-specific gap analysis, a question about training needs will be included in the quality survey sent out after EXCELLERAT training events, similarly to what is already done by other CoEs, such as HiDALGO.

4.2 Performed EXCELLERAT Training Activities

As described in D5.1, training activities within EXCELLERAT must aim at filling specific gaps in the present offer of the involved partners. On the one hand, the audience demand and the results of the training assessment are considered, e.g. by soliciting training of type "Data in HPC" whenever an EXCELLERAT tool is mature enough, or if an Interest Group can deliver such a training. On the other hand, EXCELLERAT provides training related to the flagship codes and tools applied to their respective technical and scientific domains.

Since D5.3, the Portal has evolved in a platform where training can be managed in a more comprehensive way compared to the Website. D4.2 *"Report on the Service Portfolio"* contains the Portal guidelines, according to which training services should be accessed from the three dashboards as one-on-one consulting services (i.e., a community addressing a specific provider). These guidelines have been implemented and the Portal acts now as the CoE Knowledge Hub in at least three ways, namely as a browser for training and education events, as a repository for learning material, and as a mean to submit requests for dedicated training.

For these reasons, the complete training area has been replaced in the EXCELLERAT Website by a link to the Portal, while the Website represents the main channel for dissemination and outreach. In the Portal, training can be intuitively related to other areas of the project (partners, application software, HPC resources, etc.) by the tag system Furthermore, training events can be displayed in a more structured way compared to the Website, by specifying event categories such as Date and Place, Presenters, URL, Link to streaming, etc. Finally, self-learning and complementary material to each event (e.g. tutorials and best-practice guidelines) can be accessed directly through each event's page.

On the right-hand side of the event-browsing page (Figure 16), three buttons allow the user to submit a request for a special training, to browse the consortium partners' own training offer also beyond EXCELLERAT, and to access the Euro-HPC training registry. This is in particular the result of a collaboration with FocusCoE. In this embedded calendar, courses by several HPC CoEs as well as by PRACE are displayed and promoted [28].
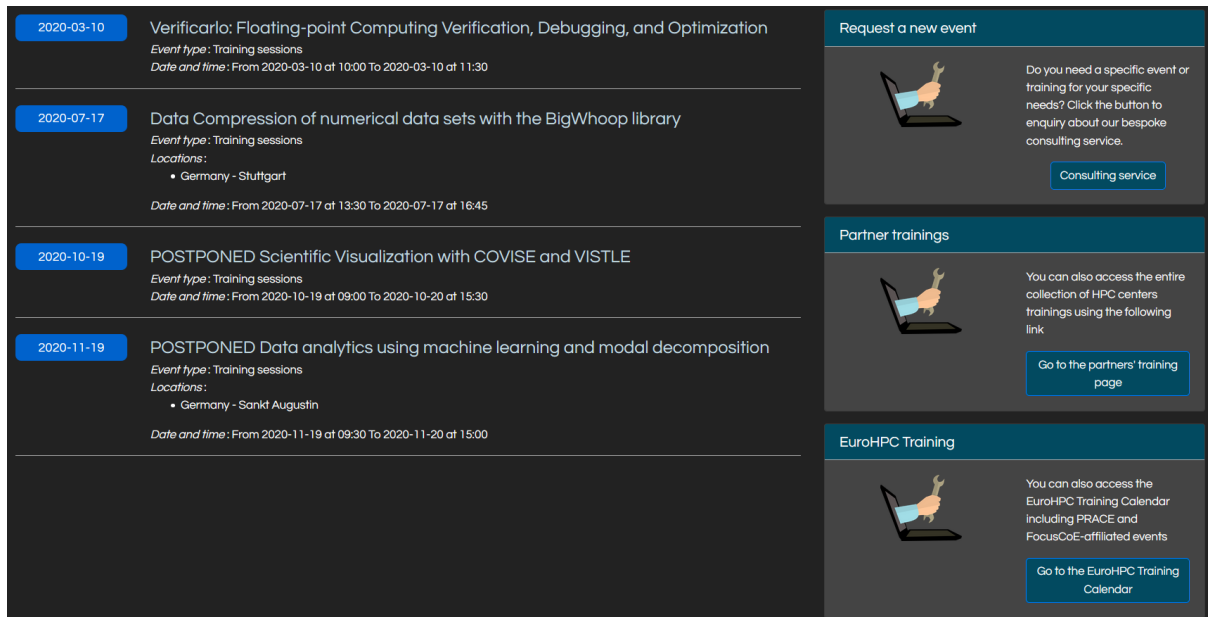
Figure 16: Training page on the Portal [29].

When listing the conducted and planned training effort, we also indicate the respective training category according to the definition introduced in D5.1 (Section 5.2), and here summarised:

- Parallel Programming [PAR]
- Computational Fluid Dynamics [CFD]
- Scientific Visualization [VIS]
- Compute Cluster: Usage and administration [CLU]
- Performance Optimization and Debugging [PRF]
- Data in HPC [DAT]
- Programming Languages for Scientific Computing [LNG]
- Scientific Conferences and Workshops [C+W]
- Training for special communities [COM]
- Others [OTH]

In M1-13 the training effort below has been conducted (as by D5.3):
- "Introductory Webinar on OpenProject" [OTH] (2019-08-02): a web-seminar on an internal service, hosted by HLRS. **Collaboration**: Seminar jointly conducted by one member of EXCELLERAT and one of HiDALGO. **Material**: A link to the recording is available on the Portal.
- "SWAN webinar" [DAT] (2019-08-26): a web-seminar on the system for the worldwide exchange of engineering data "SWAN" developed at SSC-Services GmbH. SWAN is the predecessor of the EXCELLERAT "Data Management & Data Transfer Platform", which allows for safe, fast, traceable, bidirectional and online data transfer among the HPC centres members of EXCELLERAT. **Material**: Links to the recording and documentation are available on the Portal**.**
- "Nek5000 tutorial and hands-on" [CFD] (2019-12-06): a use-case presentation and hands-on session by KTH personnel on Nek5000 at the "School on Numerical Methods for Parallel CFD" at CINECA in Rome. **Collaboration**: Two additional project (ChEESE CoE [30] and HiFi-TURB [31]) were presented at this CFD School. **Material** is available on the Portal.
- "BigWhoop compression library" [DAT] (2019-11-08): a short presentation and demonstration by an USTUTT developer within the first EXCELLERAT Interest Group

meeting. BigWhoop is a compression tool for effectively reducing the size of large numerical data sets, being developed as part of the EXCELLERAT "Data Management & Data Transfer Platform" for efficient data transfer among HPC systems. **Material**: Recording is available on request.

In M14-23 the training effort below has been conducted:

- "Verificarlo: Floating-point Computing Verification, Debugging, and Optimization" [PRF] (2020-03-10) This web-seminar organised by TERATEC and conducted by Eric Petit aimed at presenting a tool external to EXCELLERAT to the consortium members. Verificarlo addresses floating point verification and debugging to support industrial HPC codes in achieving reproducibility, portability across hardware and software, and performance optimisation. **Material**: Recording is available on request. The Portal contains a link to the Verificarlo github repository and PDF material.

- "Data compression of numerical data sets with the BigWhoop library" [DAT] (2020-07-17): This introduction and hands-on session was part of the four-day Deep Learning Workshop at HLRS, consisting of a three-day NVIDIA GPU workshop followed by one practical day by USTUTT personnel. BigWhoop is a solution proposed by EXCELLERAT to circumvent the I/O bottleneck generated by data from engineering computing and Machine Learning applications. The presentation at this workshop aimed at giving a quick overview of the BigWhoop compression library, its application to numerical datasets, and how its lossy compression stage affects the datasets information content (Figure 17). **Material** is available on the Portal.
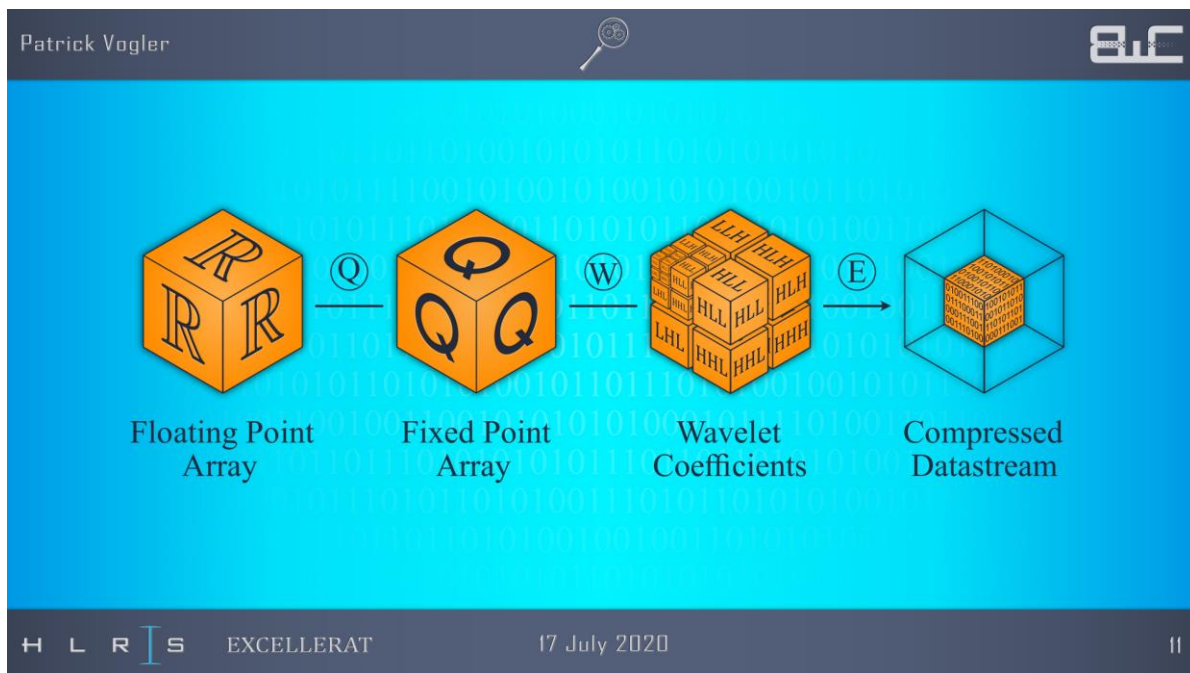

Figure 17: Data Compression with BigWhoop [32].

Both the "Nek5000" and the "Data compression" workshops are examples of EXCELLERAT-specific complementary interdisciplinary aspects which integrate activities of the consortium members with a broader target (a CFD school and a Machine Learning course respectively). Further "integration" activities by EXCELLERAT are planned and can be found in the next section.

The other initiatives aimed at instructing on the use of tools (e.g. OpenProject) used within EXCELLERAT and other CoEs, or at filling gaps in the consortium training offer with new

topics (e.g. floating-point verification with Verificarlo, and in-situ visualisation with Vistle, in the next section).

Further events were planned in M14-23 but due to the COVID-19 pandemic outbreak had to be rescheduled because of difficulties in re-organising and switching to an online format at short notice. Activities that had to been postponed have been: "Scientific Visualization with COVISE and VISTLE", "Data analytics using machine learning and modal decomposition", and the "AMD Machine Learning workshop". Details on these and other future events are provided in the next section.

4.3 Planned Training and Workshop Activities

In D5.3, the importance of in-class events as an efficient mean of knowledge transfer was highlighted, in particular in the perspective of integrating EXCELLERAT-specific training sessions within the partners' established curricula of courses and workshops. After the COVID-19 pandemic, this strategy had to be revised since in-class training had to be completely replaced by online training. The conversion was mostly carried out surprisingly fast, thanks to a prompt response of the partners in developing new or getting familiar with existing online tools. Collaboration among the CoEs also positively contributed, see e.g. the FocusCoE and E-CAM "Guide for capturing training events" [33]. Even so, many established courses could only partially or could not be converted to online ones. Therefore, EXCELLERAT had to proceed somehow carefully with the "integration" strategy.

Nevertheless, the "Data Compression"-BigWhoop presentation and hands-on session (previous section) has been a successful example of EXCELLERAT contribution in an existing course. This will be followed by these integrated events:

- "Scientific Visualization with COVISE and Vistle" [VIS]: The USTUTT visualisation department is also developing within EXCELLERAT the tool Vistle, an extensible software environment that integrates simulations on supercomputers, post-processing and parallel interactive visualisation in immersive virtual environments. A combined COVISE-Vistle course is planned for EXCELLERAT and external (industrial or post-graduate) participants in 2021. The two-day COVISE visualisation course is normally offered at USTUTT twice a year, but in 2020 it has not been possible to offer this course remotely. The USTUTT visualisation cluster cannot host a high number of remote accesses, and at the same time Vistle installers for local systems are a work in progress.
- "FRAUNHOFER Data Analysis" [DAT]: In Y3, the EXCELLERAT partner FRAUNHOFER plans to extend their current training on "Data Scientist" and "Big Data Architecture" within the "Fraunhofer Academy" with training modules on Data Analysis and Engineering Simulation Data for industrial end-users, as soon as the technical application reaches maturity (see also D4.2). Such workshops would be further instances of the "Data Analytics" workshop described below, with an enhanced application focus and targeted to a greater public.

Further planned events are listed below:

- "Data analytics using machine learning and modal decomposition" [DAT]: The workshop will be structured into two parts. The first part provides a general introduction to Machine Learning (ML) aiming at the analysis and post-processing of engineering data extracted from simulations. In agreement with the participants, supervised and unsupervised ML techniques are discussed with application in clustering, modelling and predicting as well as visualising intrinsic structures of the data in a low dimensional space. The second part of the workshop aims at instructing how to perform a modal decomposition of unsteady three-dimensional simulation data using dynamic mode

decomposition (DMD) and proper orthogonal decomposition (POD) on the EXCELLERAT software tool for Modal Decompositions developed at RWTH (see Section 4.3 of D4.6).

This workshop targets internal members of EXCELLERAT and Interest Groups representatives. The lecturers (FRAUNHOFER and RWTH personnel) are currently elaborating a consortium- and Interest Groups-internal survey to define the workshop content, and evaluating strategies to provide this course online in Q1/2021. Especially the hands-on part which foresees work on the participants' own data would pose some obstacles and might be postponed to a live event later on. This activity is under preparation at the time of writing (Figure 18).

- "AMD Machine Learning workshop" [COM] (January 2021, to be confirmed): In the framework of a collaboration between AMD and USTUTT, the EXCELLERAT Interest Group AMD will provide their first training of this kind in Europe, properly embedded into the USTUTT training activities. This remote training will allow participants to actively use the AMD GPUs through the open source development platform ROCm for HPC GPU computing [34]. Initial plans include a Machine Learning teaching module, and an HPC module on profiling and debugging tools for ROCm. Due to COVID-19 and delays in the infrastructural deployment, this activity had to be postponed, even though planned already since Q1/2020. This activity is under preparation at the time of writing.

- "First Joint CoE Technical Workshop" [C+W] (2021-01-27 to 29): Initially planned as a follow-up of the "CFD School" (Section 4.1), the organising team by members of USTUTT, TERATEC and CINECA, eventually opted for a more CoE-inclusive event. It will consist of four technical sessions (load balancing, in-situ and remote visualisation, co-design, GPU porting), each with presentations by the three CoEs ChEESE, HiDALGO, and EXCELLERAT (and a few guest presentations). Participation to each session will be limited to maximise the exchange and the mutual learning effect. This three-day workshop had been planned with large advance as a post-pandemic joint effort of many CoEs at USTUTT, but had eventually to be converted to an online event as well.

- "Nek5000 and FEniCS" [CFD]: Long-term plans include also another training on the flagship applications NEK5000 and the Finite Element solver FEniCS. Different aspects of FEniCS could be taught in parallel by both EXCELLERAT and external experts: C++ HPC FEniCS (by KTH and CINECA with focus on CFD), and python FEniCS for Model Order Reduction (by contacts at SISSA mathLab, Trieste, Italy).
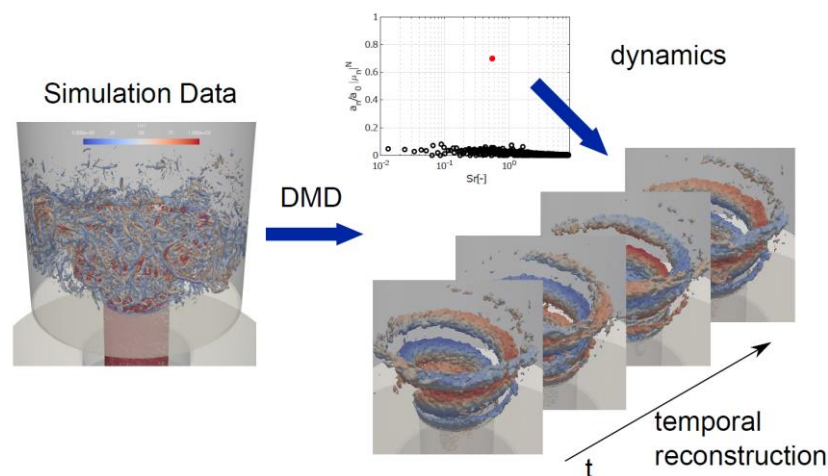

Figure 18: Data Analytics workshop [35], image courtesy Institute of Aerodynamics RWTH Aachen.

4.4 KPI Analysis

KPI number	Description	Definition	Target	Score at M13	Target M15-M23	Score at M23
5.4.1	Number of participants in EXCELLERAT training (per year)	*N_participants_per_year*	**≥20** for the first year after MS4 completion, +20% for the final period	72	15	79
5.4.2	Size of developed material in number of lecture units	*Size_training_material*	**≥30** at the end of the first year after MS4 completion, +50% at the project conclusion	6	22.5	10
5.4.3	Gender ratio of the participants	*R_female_male_quotient*	50%	14%	50%	24%
5.4.4	ACSI score: Satisfaction about EXCELLERAT training outcome (via survey/feedback sheet) as average per year	$N_{ACSI}\_training\_per\_year$	≥80	71	≥80	79
5.4.5	ACSI score: Quality and uptake of developed EXCELLERAT material also by non-EXCELLERAT stakeholders as average per year (via survey/feedback sheet)	$N_{ACSI}\_G\_quality\_material$	≥80	66	≥80	86

Table 5: KPIs for Task 5.4

In D5.1, five functional and quality-related KPIs have been proposed, which can be seen in Table 5. The quality-related KPIs refer to the ACSI Score (American Customer Satisfaction Index) defined in D5.1 (Section 2.3). In the Table, the global targets as defined in D5.1 are reported, as well as the scores at M13 (as in D5.3), the relevant target for this deliverable, and the current scores at M23. For 5.4.1, 5.4.3, 5.4.4 and 5.4.5, the average score per year after MS4 (first Portal release) must be computed: For this reason, the target for 5.4.1 and 5.4.2 has been rescaled to a period of nine months (M15 to M23).

KPI 5.4.1 ("Number of participants") is again largely satisfactory. Online events, especially when integrated in partners' activities, can profit of a very large audience. It was the case e.g. of the "Data Compression" workshop at USTUTT. Web-seminars as stand-alone EXCELLERAT events can be promoted through the FocusCoE network of CoEs, in case their content can be shared outside the consortium (e.g. the "First Joint CoE Technical Workshop").

The training service of the Portal, now also acting as the CoE Knowledge Hub, allows, as foreseen in D5.3, to obtain a comprehensive and structured overview of the developed training material and to get an exact measurement of KPI 5.4.2. Many training events are expected in the next months (Section 4.3), after the Q2 and Q3/2020 decrease as a pandemic effect, and consequently an upswing in the number of material units. At the same time, most flagship applications and tools will reach maturity during Y3, and tutorials will be offered.

The goal of ≥50% not-male participants could still not be reached, even though the period M15-M23 marks an improvement compared to the D5.3 result. It is still an arguably ambitious target for training in engineering and technical domains. Even so, we expect further improvements through a closer collaboration with WP7 and the FocusCoE network in promoting the events. The CSA for the National Competence Centres (CASTIEL), started in September 2020) will also be taken in the loop to distribute the EXCELLERAT activities to the diverse European countries and stakeholders.

Compared to the results in D5.3, the targets for both KPIs 5.4.4 and 5.4.5 could be reached (or very narrowly missed) for the period M15-M23. As by D5.3 plans, an EXCELLERAT-specific online survey has been elaborated to evaluate the project's events, and to also inquire on training needs and wishes. The tool easyfeedback [36] guarantees the European storage of the collected data. EXCELLERAT contributions to external activities might still have to rely on other surveys, for example (customisable) PRACE-questionnaires.

5 HPC Provisioning

Task 5.5 "HPC Provisioning" implements the CoE hub for internal HPC resources provisioning, to support other WPs' activities. The task compiles and updates the list of available services, provides a point of contact for all CoE members to access specific services, support and documentation, and makes sure that the requests coming from the CoE are addressed by the service provisioning functions of the HPC centres. In other words, this task implements a link between the CoE and the European HPC infrastructure, including both PRACE members and EuroHPC pre-exascale prototype owners.

5.1 Provisioned Resources and Allocation

Computational resources have been provided to the Consortium, specifically by PRACE [37]. 0.5% of the total resources available for each PRACE call are reserved for the CoEs as selected by the European Commission under the E-INFRA-5-2015 and INFRAEDI-02-2018 calls for proposals. EXCELLERAT asked for the computational resources available for the Call 18, 19, 20 and 21. The assigned resources are reported in Table 1 below. For details regarding the architecture's specifications of the involved HPC systems, please see Table 3 in Annex 1: Internal Resources.

Awards (cores/hours)	PRACE 18	PRACE 19	PRACE 20	PRACE 21	Total
Marconi BDW	80,000	45,000			125,000
Marconi KNL	1,000,000	750,000			1,750,000
Marconi100			1,620,000	875,000	2,495,000
HAWK		1,150,000	1,000,000	550,000	2,700,000
JUWELS Cluster	100,000	175,000	70,000	45,000	390,000
JUWELS Booster				55,400	55,400
Joliot Curie Rome		1,715,000	950,000	350,000	3,050,000
Joliot Curie KNL		150,000	1,120,000	240,000	1,510,000
Joliot Curie SKL	350,000	180,000	170,000	170,000	870,000
MareNostrum4	700,000	240,000	150,000	240,000	1,330,000
SuperMUC-NG			302,500	86,000	388,500
Piz Daint	450,000	850,000	935,000	250,000	2,485,000

Table 6: HPC Resources allocated to Call PRACE 18-21 (cores-hours)

The accounts are open on whole the different clusters and the access is granted on request to the various partners.

The wiki page of EXCELLERAT has been updated with all the up-to-date information regarding the HPC access[2]. Specifically, a new section has been added, *HPC Provisioning*, where HPC resources available to the Consortium provided by PRACE are updated, see screenshot below:

---

[2] https://kb.hlrs.de/excellerat/index.php/EXCELLERAT_Wiki

== HPC Provisioning ==
- HPC resources on the various cluster for the Consortium provided by PRACE (Task 5.5)
- Last update: 09 December 2020
- The following page summarizes the different HPC Access procedures

Awards (Allocated)	Galileo Broadwell (CINECA): core/hours	Marcon100 (CINECA): local core/hours-(node/hours)	Hawk (HLRS)	JUWELS (Julich) core/hours	Joliot Curie AMD (Rome)	Joliot Curie KNL	Joliot Curie Irene SKL	MareNostrum4 (BSC)	Piz Daint (CSCS)-(node-hours)	SuperMUC-NG
PRACE Call 18	80.000	50.000 - (1.562)	--	--	--	--	Cumulative with other Calls	--	--	
resource used	148.860	0	--	--	--	--	--	--	--	
resource used (Percentange %)	**186%**	0%	--	--	--	--	--	--	--	
Start / End allocation	19-08-2019 / 18-08-2020	10-07-2019 / 08-08-2020	--	--	--	--	--	--	--	
Account	Pra18_Excell_0	Pra18_Excelera_0	--	--	--	--	--	pr1ejz00	prce10	
PRACE Call 19	45.000	37.500 - (1.171)	1.150.000	--	150.000	180.000	Cumulative with other Calls		--	
resource used	213.546	5.660	--		--	150.000	180.000		--	
resource used (Percentange %)	**475%**	15%	--	--	--	--	--		--	
Start / End allocation	30-10-2019 / 06-01-2021	30-10-2019 / 06-01-2021	--	--	--	--	--			
Account	Pra19_Exceller_0	Pra19_Exceller_1	--	--	RACOE006	RACOE006	RACOE006	pr1ejz00	prce10	
PRACE Call 20	--	147.272 (4602)	1.000.000	275.000	Cumulative with Call 21	Cumulative with Call 21	Cumulative with Call 21	Cumulative with other Calls		
resource used	--	0	0	0	--	--	--	--	--	
resource used (Percentange %)	--	0%	0%	0%	--	--	--	--	--	
Start / End allocation	--	24-04-2020 / 26-04-2021	/ 31/10/2020	01-07-2019 / 31-03-2021	--	--	--		--	
Account	--	Pra20_Exceller	--	prcoe05	--	--	pr1ejz00		prce10	
PRACE Call 21	--	79.546	550.000	45.000	1.255.138	240.000	170.000	540.000	1840 (Cumulative with previous Calls)	86.000
resource used	--	0	0	0	272943	0	0	197.77	13 (to be used on quartelry basis)	-
resource used (Percentange %)	--	0%	0%	0%	22%	0	0	37%	1%	-
Start / End allocation	--	08-10-2020 / 31-03-2021	/ 31/10/2020	01-07-19/30-09-2021	16-10-2020/01-04-2021	--			01-07-2019/01-05-2021	-
Account	--	Pra21_EXCELLERAT	--	prcoe05	RACOE006	--	--		prce10	
Upon Request/Special Petition	--	--	--	--	--	--	--	--	--	
resource used	--	--	--	--	--	--	--	--	--	
resource used (Percentange)	--	--	--	--	--	--	--	--	--	
Start / End allocation	--	--	--	--	--	--	--	--	--	
Account	--	--	--	--	--	--	--	--	--	

Figure 19: Screenshot of HPC Provisioning from the wiki

Moreover, a further section "HPC Access"[3] has been added on the different procedure for HPC access on the various systems,

Figure 20: Screenshot of HPC Access

A useful guide has been recently added on the same page on how to authenticate with multiple ssh keys on the various HPC system, to be compliant with the every-day stringent access policies.

Figure 21: Screenshot of Authentication with multiple ssh keys from "HPC Access" page

5.2 KPI Analysis

Table 7 presents the KPIs for this task, as defined in D5.1.

In particular, we can mention one project that has been approved. STRIKE Project, PI Cerfacs. Budget requested: 30 M core/hours. Moved to Irene KNL, with budget reduction.

KPI number	Description	Definition	Target	Score at M23
5.5.1	Percentage of approved proposals in competitive calls for resources for EXCELLERAT projects	N_prj_appr / $N\_prj\_submitted$ * 100	>50%	100

---

[3] https://kb.hlrs.de/excellerat/index.php/HPC_Access

| 5.5.2 | Number of yearly allocated cpu_hours for EXCELLERAT projects | N_cpu_hours | $\geq 10^6$ for PY2, $\geq 10^7$ for PY3 | ca. 65 M |

Table 7: KPIs for Task 5.5

6 Conclusion and Time Plan

With the successful submission of this deliverable D5.5 and its companion D5.4 at M25, last of a long series of preparatory deliverable reports, Milestone MS4 "*First verification and update phase of the EXCELLERAT services has been performed*" is achieved.

In year 3 the work will continue in all tasks. This means in particular that the service portal will evolve according the feedback received by partners and users, inserting a new dashboard for vendors and achieving a better integration with the Website. The service portfolio will be expanded, implementing the directives obtained by the other work packages, and a significant effort will be devoted in collaboration with WP7 to expand the external user base.

Furthermore, while the activities for the new selected applications will be completed leading to success stories, new onboarding opportunities will be assessed. Similarly, the training portfolio will be expanded and will evolve adapting to the new post-COVID reality.

Finally, Task 5.5 will continue to provide the necessary resources for WP2, 3 and 4 activities.

7 References

[1] "BSCW website," [Online]. Available: https://www.bscw.de/en/. [Accessed 17 12 2019].

[2] "Nextcloud website," [Online]. Available: https://nextcloud.com.

[3] "Mediawiki website," [Online]. Available: https://www.mediawiki.org/wiki/MediaWiki.

[4] "Etherpad website," [Online]. Available: https://etherpad.org/. [Accessed 17 12 2019].

[5] "GoToMeeting website," [Online]. Available: https://www.gotomeeting.com/. [Accessed 17 12 2019].

[6] "OpenProject website," [Online]. Available: https://www.openproject.org/. [Accessed 17 12 2019].

[7] "Interest Groups in the EXCELLERAT Portal," [Online]. Available: https://services.excellerat.eu/searchpartners/interestgroup.

[8] "Interest Groups in the EXCELLERAT Portal," [Online]. Available: https://services.excellerat.eu/searchpartners/interestgroup.

[9] F.Battista, P.Gualtieri, J.-P.Mollicone and C.M.Casciola, "Application of the Exact Regularized Point Particle method (ERPP) to particle laden turbulent shear flows in the two-way coupling regime," *International Journal of Multiphase Flow,* vol. 101, pp. 113-124, 2018.

[10] "ROSITA (ROtorcraft Software ITAly)," [Online]. Available: https://home.aero.polimi.it/flowmesh/description_d2.html.

[11] J. Steger, F. Dougherty and J. Benek, "A Chimera Grid Scheme, Advances in Grid Generation,," in *Applied Mechanics, Bioengineering, and Fluids Engineering Conference*, Houston, TX, 1983.

[12] "HiDALGO CoE website," [Online]. Available: https://hidalgo-project.eu/. [Accessed 17 12 2019].

[13] "HiDALGO Collaborations," https://hidalgo-wiki.hlrs.de/en/tasks/collaborations.

[14] "HPC and Big Data Technologies for Global Systems: Interactive Workshop and Hands-on Session," [Online]. Available: https://www.hipeac.net/2020/bologna/#/program/sessions/7722/.

[15] "UAP and EXCELLERAT slidedeck," [Online]. Available: https://drive.google.com/file/d/1AdIgqtabD6GAAWLiCBlefwLbXjMfL7d4/view.

[16] "Onboarded application software in the EXCELLERAT Portal," [Online]. Available: https://services.excellerat.eu/searchcodes/onboarded.

[17] "Bizagi modeler website," [Online]. Available: https://www.bizagi.com/en/products/bpm-suite/modeler. [Accessed 17 12 2019].

[18] "Service requests feature in the EXCELLERAT Portal," [Online]. Available: https://services.excellerat.eu/searchrespsis.

[19] "Onboarded application software in the EXCELLERAT Portal," [Online]. Available: https://services.excellerat.eu/searchcodes/onboarded.

[20] "Survey to measure the satisfaction of the onboarding process for a service," [Online]. Available: https://easy-feedback.de/preview/1244159/wwT14a-700f01dba88c9739a79762926a3c6fb2.

[21] "The finite volume method for solid mechanics (project)," [Online]. Available: https://www.researchgate.net/project/The-finite-volume-method-for-solid-mechanics.

[22] A. Heft, T. Indinger and N. Adams, "Experimental and Numerical Investigation of the DrivAer Model," in *ASME 2012*, Puerto Rico, USA, 2012.

[23] "1st Automotive CFD Prediction Workshop," [Online]. Available: https://autocfd.eng.ox.ac.uk/.

[24] "Survey to measure the satisfaction of the onboarding process for a service," [Online]. Available: https://easy-feedback.de/preview/1244159/wwT14a-700f01dba88c9739a79762926a3c6fb2.

[25] "Survey to measure the satisfaction of the processed applicants," [Online]. Available: https://easy-feedback.de/preview/1244169/58Qkv7-4c9a22a95b0f02b0e3a48ef79627cef0.

[26] "CASTIEL project website," [Online]. Available: https://www.castiel-project.eu/.

[27] "FocusCoE CSA website," [Online]. Available: https://www.focus-coe.eu/.

[28] "Pan-European Training Registry," [Online]. Available: https://services.excellerat.eu/ftrainings.

[29] "Training page on the Portal," [Online]. Available: https://services.excellerat.eu/searchevents/training.

[30] "ChEESE CoE website," [Online]. Available: https://cheese-coe.eu/.

[31] "HiFi-TURB project," [Online]. Available: https://www.ercoftac.org/hifi-turb-project/.

[32] "Material for "Data Compression with BigWhoop"," [Online]. Available: https://services.excellerat.eu/api//excellerat/filedownload_event/11/DeepLearningWorkshop2020-BigWhoop___POINT___pdf.

[33] "Guide for capturing training events," [Online]. Available: https://training.e-cam2020.eu/datasets/5e8b24e6e4b03c89fa46f941.

[34] "The AMD ROCm Platform," [Online]. Available: https://rocmdocs.amd.com/en/latest/index.html.

[35] "Data analytics using machine learning and modal decomposition," [Online]. Available: https://services.excellerat.eu/viewevent/8.

[36] "easyfeedback website," [Online]. Available: https://easy-feedback.com/.

[37] "PRACE RI website," [Online]. Available: http://www.prace-ri.eu/. [Accessed 17 12 2019].

8 Annex 1: EXCELLERAT Onboarding Proposal: PAinG-Flow

Organisation (team)
FLUMACS

Members
Prof. C. M. Casciola - University of Rome Sapienza, Mechanical and Aerospace Department
Prof. P. Gualtieri - University of Rome Sapienza, Mechanical and Aerospace Department
Prof. F. Battista - University of Rome Sapienza, Mechanical and Aerospace Department

Description of work
The project purpose is the implementation of the Exact Regularized Point Particle (ERPPR) method to exert the particle/fluid momentum coupling in multiphase turbulent flows on a code able to run on supercomputers with GPUs. The code will solve the incompressible Navier-Stokes equation in wall-bounded and free-shear conditions. The fluid Eulerian phase is solved by discretizing the Navier-Stokes equation on staggered grid with a centred second-order finite difference scheme in space and with a third-order four-steps Low-Storage Runge-Kutta scheme in time. The solid phase is modelled by evolving each single particle in a Lagrangian way with the same temporal scheme of the Eulerian phase.

The method aims at regularizing the force that a small particle exerts on the fluid on a physical ground. As a particle moves along its trajectory, it generates a small-scale vorticity field that can be evaluated in a closed form. Indeed, the disturbance flow is shown to obey the unsteady Stokes problem with appropriate boundary conditions to account for the wall effects. The disturbance field produced on the scale of the particle is then regularized by the fluid viscosity. Hence, the process of vorticity generation and viscous diffusion allows to model on a physical ground without any "ad hoc" numerical artefact the inter-phase momentum coupling. The code is written in Fortran 90 language and uses MPI library for parallelization.

Workplan
- **2 months**: study of the implementation algorithms and of the GPU technologies suitable for the implementation
- **4 months**: implementation and porting of the code on GPU machine with particular care to the evolution of the Lagrangian phase.
- **2 months**: testing of the new implementation and first production run.
- **1 month**: study of the possibility to implement several geometries in the code and production of the project report.

Computing machine
Cineca Marconi100 cluster
Requested computing time: 10000 node hour

Requested EXCELLERAT support
Cineca HPC specialist: F. Salvadore

9 Annex 2: Assessment of Training Activities 2020

Category:	Parallel Programming	CFD	Performance Opt. & Debugging	Data in HPC	Programming Languages	Scientific Visualization	Compute Cluster – Usage & Adm.	Training for Special Communities	Conferences and Scientific Workshops
# of events 2019	41	20	14	13	20	4	6	11	9
# of events 2020	32	18	11	19	28	4	5	9	8
Trend	**-22%**	-10%	**-21.4%**	**+46.2%**	**+40%**	=	-16.7%	-18.2%	-11.1%

Figure 22: Training activities 2020 assessment

10 Annex 3: Internal Resources

The resources available to the Consortium are continuously updated. The up-date list of Tier-0 resources is provided by the most recent Prace Call. The Table reported below is extracted from the Technical Guidelines of the Prace Call 22[4].

		HAWK	Joliot-Curie Rome	Joliot-Curie SKL	Joliot-Curie KNL	JUWELS Cluster	JUWELS Booster	Marconi100	Mare Nostrum 4	Piz Daint	SuperMUC-NG
	System Type	HPE	Bull Sequana	Bull Sequana	Bull Sequana	Bull Sequana	Bull Sequana	IBM Power 9 AC922 Whiterspoon	Lenovo	Hybrid Cray xC50	Lenovo ThinkSystem
Compute	Processor type	AMD Epyc Rome	AMD Epyc Rome	Intel Xeon Platinum 8168 2.7 GHz	Intel Knights Landing	Intel Xeon Skylake Platinum 8168	AMD EPYC Rome	2 * IBM POWER9 AC922 at 3.1 GHz per node	Intel Xeon Platinum 8160 2.1 GHz	Intel® Xeon® E5-2690 v3 @ 2.60GHz (12 cores)	Intel Skylake Xeon Platinum 8174
	Total nb of nodes	5 632	2 292	1 656	828	2 511	T.B.A.	980	3 456	5 704	6480
	Total nb of cores	720 896	293 376	79 488	52 992	120 528	T.B.A.	31 360	165 888	68 448	31 1040
	Nb of accelerators /node	n.a.	n.a	n.a.	n.a.	n.a.	4	4 GPU per node	n.a.	1 GPU per node	n.a.
	Type of accelerator	n.a.	n.a	n.a.	n.a.	n.a.	NVIDIA next Gen.	NVIDIA® Volta® V100, Nvlink 2.0, 16GB	n.a.	NVIDIA® Tesla® P100 16GB	n.a.
Memory	Memory / Node	256 GB	256 GB	192 GB	96 GB DDR4 + 16 GB MCDRAM	96 GB	T.B.A.	256 GB DDR4 + up to 1.6 TB Optane Memory per node	96 GB (200 nodes with 384GB)	64 GB	96 GB

---

[4] https://prace-ri.eu/wp-content/uploads/Technical_Guidelines_Call_22.pdf

Network	Network Type	Infiniband HDR	Infiniband HDR 100	Infiniband EDR	BULL BXI	InfiniBand EDR	InfiniBand HDR	Mellanox Infiniband EDR	Intel Omni-Path Architecture	Cray Aries	Intel Omni-Path Architecture
	Connectivity	9D enhanced Hypercube	Dragonfly+	Fat Tree	Fat Tree	Fat Tree	Dragonfly+	DragonFly+	Fat Tree	Dragonfly	Fat tree within island (786 nodes) pruned tree between islands

		HAWK	Joliot-Curie	JUWELS	Marconi100	MareNostrum 4	Piz Daint	SuperMUC-NG
Home file system	type	NFS	NFS	GPFS	GPFS	GPFS	GPFS	GPFS
	capacity	100 TB	0.5 TB	2.8 TB	200 TB	32 TB	160 TB	256 TB
Work file system	type	Lustre	Lustre	GPFS	GPFS	GPFS	GPFS	GPFS
	capacity	25 PB	9.2 PB	2.3 PB	3 PB	4.3 PB	6.3 PB	33 PB
Scratch file system	type	n.a.	Lustre	GPFS	GPFS	GPFS	Lustre	GPFS
	capacity	n.a.	5.2 PB	9.1 PB	2 PB	8.7 PB	8.8 PB	17 PB
Archive	capacity	On demand	On demand	On demand	On demand	n.a.	n.a.	On demand
Minimum required job size	Nb of cores	8192	1 024		2 nodes	1024	6 nodes	960

Table 8: Updated list of internal available HPC resources